

Importance Sampling, Large Deviations, and Differential Games

Paul Dupuis* and Hui Wang†

Lefschetz Center for Dynamical Systems
Brown University
Providence, R.I. 02912
USA

Abstract

A heuristic that has emerged in the area of importance sampling is that the changes of measure used to prove large deviation lower bounds give good performance when used for importance sampling. Recent work, however, has suggested that the heuristic is incorrect in many situations. The perspective put forth in the present paper is that large deviation theory suggests many changes of measure, and that not all are suitable for importance sampling. In the setting of Cramér's Theorem, the traditional interpretation of the heuristic suggests a fixed change of distribution on the underlying independent and identically distributed summands. In contrast, we consider importance sampling schemes where the exponential change of measure is adaptive, in the sense that it depends on the historical empirical mean. The existence of asymptotically optimal schemes within this class is demonstrated. The result indicates that an adaptive change of measure, rather than a static change of measure, is what the large deviations analysis truly suggests. The proofs utilize a control-theoretic approach to large deviations, which naturally leads to the construction of asymptotically optimal adaptive schemes in terms of a limit Bellman equation. Numerical examples contrasting the adaptive and standard schemes are presented, as well as an interpretation of their different performances in terms of differential games.

*Research of this author supported in part by the National Science Foundation (NSF-DMS-0072004, NSF-ECS-9979250) and the Army Research Office (DAAD19-00-1-0549, DAAD19-02-1-0425).

†Research of this author supported in part by the National Science Foundation (NSF-DMS-0103669).

1 Introduction and Background

A basic technique for the approximation of probabilities and functionals of probability measures is Monte Carlo simulation. Let X be a random variable taking values in the real numbers. To estimate EX , a sequence of independent and identically distributed (iid) copies X_0, X_1, \dots of X are generated, and the estimate for EX based on the first K samples is just the sample mean: $Q_K \doteq (X_0 + X_1 + \dots + X_{K-1})/K$. Since the convergence of the estimate is based on the law of large numbers, a standard rate of convergence can be defined in terms of the variance $\text{var}[X_0]$. Indeed, the variance of Q_K (assuming that $\text{var}[X_0]$ exists) is just $\text{var}[X_0]/K$.

In cases where the variance is large, and especially if it is large compared to EX , it may take a large number of samples before the variance of the estimator is acceptably small relative to EX . This is a common occurrence when estimating rare events, and also when estimating functionals whose behavior is largely determined by rare events. In such cases, one may be tempted to use some form of *importance sampling* to reduce the variance, and hence speed up the computation by requiring fewer samples. Since importance sampling is most effective when dealing with rare events, it is perhaps not surprising the literature on importance sampling and its relation to large deviations is extensive.

The basic formulae of importance sampling are as follows. Suppose that X has distribution θ , and consider an alternative sampling distribution τ . It is required that θ be absolutely continuous with respect to τ , so that the Radon-Nikodym derivative $f(x) \doteq (d\theta/d\tau)(x)$ exists. Iid samples $\bar{X}_0, \bar{X}_1, \dots$ with distribution τ are generated, and the estimate

$$\bar{Q}_K \doteq \frac{1}{K} \sum_{k=0}^{K-1} \bar{X}_k f(\bar{X}_k)$$

is considered in lieu of Q_K . Since

$$E\bar{X}_k f(\bar{X}_k) = \int_{\mathbb{R}} x f(x) \tau(dx) = \int_{\mathbb{R}} x \theta(dx) = EX,$$

\bar{Q}_K is an unbiased estimate of EX , with a rate of convergence determined by

$$\text{var} [\bar{X}_0 f(\bar{X}_0)] = \int_{\mathbb{R}} x^2 f(x) \theta(dx) - \left[\int_{\mathbb{R}} x \theta(dx) \right]^2.$$

The optimization of this quantity over all possible τ is inappropriate. For example, suppose θ is supported on $[0, \infty)$, and $\theta(dx) = g(x) dx$. Let

$m \doteq EX = \int_0^\infty x\theta(dx)$ and $\tau(dx) \doteq m^{-1}xg(x)dx$. Then θ is absolutely continuous with respect to τ , with $f(x) = m/x$. Furthermore,

$$\text{var} [\bar{X}_0 f(\bar{X}_0)] = \int_{\mathbb{R}} x^2 f(x)\theta(dx) - m^2 = 0.$$

However, such a distribution τ is of little use in practice since it requires knowledge of m , the very thing we want to estimate! Instead of this unconstrained optimization, one typically seeks to minimize over parameterized families of alternative sampling distributions.

When the distribution of X_0 is connected to a large deviations problem, certain sampling distributions are immediately suggested by the form of the large deviations variational problem. In order to explain this connection we specialize to the setting of Cramér's Theorem, which will be used throughout the rest of the paper. However, the conclusions we will draw on the relations between importance sampling, large deviations and differential games hold in much greater generality, and some of these generalizations will be pursued elsewhere.

Let Y_0, Y_1, \dots be a sequence of iid \mathbb{R}^d -valued random variables with distribution μ , and assume the moment generating function $\int_{\mathbb{R}^d} \exp\langle \alpha, y \rangle \mu(dy)$ is finite for all $\alpha \in \mathbb{R}^d$. Let $S_n = Y_0 + \dots + Y_{n-1}$. For Borel sets $A \subset \mathbb{R}^d$, Cramér's Theorem is concerned with the large deviation approximation of $p_n \doteq P\{S_n/n \in A\}$. Let $1_A(y)$ denote the function equal to 1 if $y \in A$ and zero otherwise. Suppose that for some fixed value n we consider Monte Carlo approximation for $P\{S_n/n \in A\}$. In terms of the notation introduced previously, we have $X = 1_A(S_n/n)$, and straightforward Monte Carlo would require the generation of many independent copies of X (i.e., of S_n), say $(X_0, X_1, \dots, X_{K-1})$. The rate of convergence of the naive estimator Q_K is determined by

$$\text{var}[Q_K] = (p_n - p_n^2)/K^2.$$

If $p_n \rightarrow 0$ as $n \rightarrow \infty$, this variance approaches zero. However, the relative error is

$$\text{relative error} \doteq \frac{\text{standard deviation of } Q_K}{\text{mean of } Q_K} = \frac{\sqrt{p_n - p_n^2}}{K p_n}.$$

Since $\sqrt{p_n - p_n^2}/p_n \rightarrow \infty$, a large sample size (i.e., K) is required for the estimator to achieve a reasonable relative error bound (at least when n is large). In fact, if p_n scales according to a large deviation principle then K must grow exponentially in n if a bounded relative error is to be maintained.

However, under the assumed condition on the moment generating function S_n/n satisfies what is called a *large deviation principle* (LDP). Define the convex function

$$H(\alpha) \doteq \log \int_{\mathbb{R}^d} \exp\langle \alpha, y \rangle \mu(dy), \alpha \in \mathbb{R}^d, \quad (1.1)$$

and its Legendre transform

$$L(\beta) \doteq \sup_{\alpha \in \mathbb{R}^d} [\langle \alpha, \beta \rangle - H(\alpha)], \beta \in \mathbb{R}^d. \quad (1.2)$$

It is well known that L is a nonnegative, proper, strictly convex function [30], [13, Lemma 6.2.3], and that $L(\beta) = 0$ if and only if $\beta = \int_{\mathbb{R}^d} y \mu(dy)$. Suppose the set A has the property that

$$\inf_{\beta \in A^\circ} L(\beta) = \inf_{\beta \in \bar{A}} L(\beta), \quad (1.3)$$

where A° and \bar{A} denote the interior and closure of A , respectively. Then [39] we have the large deviation approximation

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log P \{S_n/n \in A\} = - \inf_{\beta \in A} L(\beta).$$

Now the distribution of X is rather complicated, and so it makes sense to consider the change of measure with respect to the underlying distribution μ of the Y_i instead. It is at this point that the large deviation theory suggests a specific alternative sampling distribution. To distinguish indices, we henceforth reserve k to index the k th simulation in the Monte Carlo scheme, and i and j for the index in the sum that defines S_n . In addition, k will appear as a superscript, and i and j as subscripts. Suppose that instead generating Y_i^k according to μ , we generate iid sequences $\{\bar{Y}_i^k\}$ according to a distribution ν . Suppose that we further restrict to measures ν that are related to μ by an *exponential tilt*, i.e., there is $\alpha \in \mathbb{R}^d$ such that

$$\nu(dy) = e^{\langle \alpha, y \rangle} \mu(dy) / Z(\alpha), \quad Z(\alpha) = e^{H(\alpha)} = \int_{\mathbb{R}^d} e^{\langle \alpha, y \rangle} \mu(dy).$$

We then form $\bar{S}_n^k/n = (\bar{Y}_0^k + \dots + \bar{Y}_{n-1}^k)/n$, so that $1_{\{\bar{S}_n^k/n \in A\}}$ plays the role of \bar{X}_k above. It is convenient to first express the Radon-Nikodym derivative in terms of the iid (in j) variables \bar{Y}_j^k (rather than the complicated variable \bar{S}_n^k/n), and so we arrive at the estimator

$$\frac{1}{K} \sum_{k=0}^{K-1} 1_{\{\bar{S}_n^k/n \in A\}} \prod_{j=0}^{n-1} \left[e^{-\langle \alpha, \bar{Y}_j^k \rangle + H(\alpha)} \right] = \frac{1}{K} \sum_{k=0}^{K-1} 1_{\{\bar{S}_n^k/n \in A\}} e^{n(-\langle \alpha, \bar{S}_n^k/n \rangle + H(\alpha))}.$$

It is not difficult to verify that this estimator is unbiased.

Recall that the rate of convergence of the estimator is determined by the variance of the summands

$$\text{var} \left[\mathbf{1}_{\{\bar{S}_n^k/n \in A\}} e^{n(-\langle \alpha, \bar{S}_n^k/n \rangle + H(\alpha))} \right].$$

We temporarily drop k from the notation, and observe that to minimize the variance (over α) it suffices to minimize the second moment

$$E \left[\mathbf{1}_{\{\bar{S}_n/n \in A\}} e^{2n(-\langle \alpha, \bar{S}_n/n \rangle + H(\alpha))} \right].$$

After calculating the moment generating function of ν and computing its Legendre transform, one can verify via Cramér's Theorem that \bar{S}_n/n also satisfies a LDP, but with L replaced by $\bar{L}(\beta) = L(\beta) + H(\alpha) - \langle \alpha, \beta \rangle$. A formal application of Varadhan's Theorem on the asymptotic evaluation of integrals [13, Theorem 1.3.4] then gives the large n approximation

$$\frac{1}{n} \log E \left[\mathbf{1}_{\{\bar{S}_n/n \in A\}} e^{2n(-\langle \alpha, \bar{S}_n/n \rangle + H(\alpha))} \right] \rightarrow - \inf_{\beta \in A} [2\langle \alpha, \beta \rangle - 2H(\alpha) + \bar{L}(\beta)].$$

We now substitute the expression for \bar{L} and optimize over α to obtain the min/max problem

$$\sup_{\alpha \in \mathbb{R}^d} \inf_{\beta \in A} [\langle \alpha, \beta \rangle - H(\alpha) + L(\beta)].$$

According to the previous discussion, the supremizing α should identify an asymptotically optimal sampling distribution among all changes of measure of the prescribed form. If the conditions of the min/max theorem hold (e.g., if A is convex and bounded), then the sup and inf can be permuted, and we find that

$$\begin{aligned} \sup_{\alpha \in \mathbb{R}^d} \inf_{\beta \in A} [\langle \alpha, \beta \rangle - H(\alpha) + L(\beta)] &= \inf_{\beta \in A} \left(\sup_{\alpha \in \mathbb{R}^d} [\langle \alpha, \beta \rangle - H(\alpha)] + L(\beta) \right) \\ &= \inf_{\beta \in A} 2L(\beta). \end{aligned}$$

Furthermore, if β^* is a solution to the problem $\inf_{\beta \in A} L(\beta)$, then a supremizing α^* can be identified as the conjugate dual of β^* , i.e., as a point that maximizes $[\langle \alpha, \beta^* \rangle - H(\alpha)]$ (assuming that this maximum exists).

The identification of the optimizing α brings in a further interesting connection with the theory of large deviations. In the traditional proof of Cramér's Theorem (as well as many other large deviation results), the lower

bound is proved by a change of measure argument, and the upper bound by a somewhat involved application of Chebyshev’s inequality [39]. Our interest is in the lower bound. Suppose that the rate function L has been guessed (or that it has been suggested by an existing upper bound), and that one wishes to estimate $P\{S_n/n \in A\}$ from below. Then one considers all exponential changes of measure on the underlying distribution μ which shift the mean to $\beta^* \in A$. By estimating the asymptotically dominant (as $n \rightarrow \infty$) term in the Radon-Nikodym derivative, one can find a lower bound for each such change of measure, and then optimize to obtain the tightest lower bound. Although we will not repeat the details of the proof here, it turns out that the correct change of measure is also characterized as the exponential tilt associated with a supremizing (or nearly supremizing) point in $\alpha \rightarrow [\langle \alpha, \beta^* \rangle - H(\alpha)]$.

Thus the asymptotically optimal change of measure that is used to prove the large deviation lower bound formally coincides with the distribution suggested by the preceding argument for use in importance sampling. This formal connection has been made rigorous in certain circumstances, and consequently given rise to the following heuristic: *the change of measure used to prove the large deviation lower bound should be a good (perhaps nearly optimal) distribution to use for purposes of importance sampling.* The first result of this type was given by Siegmund [37]. The basic idea was subsequently investigated in many contexts, and a small selection of the literally hundreds of papers on the topics is [1, 2, 3, 7, 8, 9, 11, 12, 14, 23, 25, 28, 29, 31, 32, 35]. Necessary and sufficient conditions under which a prescribed scheme is asymptotically optimal are discussed in [10, 33, 34], while [26] gives a survey of rare-event simulation.

However, more recent work has challenged the validity of the heuristic. For example, in [24] it is asserted that if one uses the change of measure suggested by large deviations, then in certain situations the corresponding importance sampling scheme has very poor properties. Examples are given to show that it can even perform worse than the standard Monte Carlo method! A simple example given in [24] considers Cramér’s Theorem in one dimension, with $A = (-\infty, a] \cup [b, \infty)$, and $EY_0 \in (a, b)$. Note that the set A here is not convex, and so the application of the min/max theorem above is no longer valid. Under these circumstances $\inf_{\beta \in A} L(\beta)$ must be achieved at a or b , and conditions are used to imply $0 < L(b) < L(a)$, so that the optimal change of measure in the large deviation lower bound is one that centers the mean under the new distribution on b . When one simulates under this new distribution the overwhelming majority of the sample means end close to b (for large n). There are, however, occasional “rogue” sample

means \bar{S}_n^k/n that end in the set $(\infty, a]$, and the indicator $1_{(-\infty, a]}(\bar{S}_n^k/n)$, appropriately multiplied by the Radon-Nikodym derivative, appears in the estimator. Unfortunately, it turns out that these Radon-Nikodym derivatives are very large, and in fact large enough that even though the outcome $S_n^k/n \in (-\infty, a]$ is very unlikely (exponentially small with exponent proportional to n), the term

$$1_{\{\bar{S}_n^k/n \in (-\infty, a]\}} e^{n(-\langle \alpha, \bar{S}_n^k/n \rangle + H(\alpha))}$$

dominates the variance as $n \rightarrow \infty$. Indeed, the second moment of this term alone could be (exponentially) larger than $e^{-2nL(b)}$ when $n \rightarrow \infty$. For example (see [24] for more details), take $\mu \sim N(1, 1)$ and $-a = b > 1$. We have $H(\alpha) = \alpha + \alpha^2/2$, $L(\beta) = (\beta - 1)^2/2$, $\inf_{\beta \in A} L(\beta) = L(b) < L(a)$. The change of measure will shift the mean to b , or equivalently $\alpha = \alpha^* = b - 1$. It is not difficult to verify that

$$\begin{aligned} & -\frac{1}{n} \log E \left[1_{\{\bar{S}_n^k/n \in (-\infty, a]\}} e^{2n(-\langle \alpha, \bar{S}_n^k/n \rangle + H(\alpha))} \right] \\ & \rightarrow \inf_{\beta \leq a} [\langle \alpha^*, \beta \rangle - H(\alpha^*) + L(\beta)] \\ & = -b^2 + 2b + 1. \end{aligned}$$

This term is smaller than $2L(b) = b^2 - 2b + 1$ if $b > 2$.

In spite of the examples of [24], one must examine the claims carefully before concluding that large deviations has little to say in such situations. The key issue, from our point of view, is our somewhat misleading use of the word “the” in the phrase “*the* change of measure used to prove the large deviation lower bound”. It turns out that there are *many* changes of measure that can be used to prove the lower bound, and one must consider this larger class if one hopes to identify importance sampling schemes that work well in great generality. Observe that the change of measure used in the large deviation lower bound treats all summands similarly, in that the same shift of distribution from μ to ν is used for all \bar{Y}_i^k . However, one could also consider shifts of the underlying distribution μ that dynamically adapt, in that the measure used to generate \bar{Y}_i^k could depend on the outcomes $\bar{Y}_j^k, j = 0, \dots, i - 1$. This distinction corresponds, in control terminology, to the difference between “open-loop” and “feedback” controls. The early paper [12] utilizes an adaptive importance sampling scheme to estimate certain escape probabilities for a Markov chain, and proves asymptotic optimality for a particular one-dimensional problem. However, the techniques used are not broadly applicable. The paper [14] uses adaptive importance sampling to estimate functionals of a small noise diffusion, though no proofs of

optimality are given. The distinction between these two methods of importance sampling was articulated in Sadowsky [33], where he refers to adaptive schemes as “sequential.”

There is an alternative approach to large deviations that is based on stochastic control methodology in which both upper and lower bounds are proved simultaneously. The use of stochastic control and logarithmic transforms to prove large deviation type results goes back to Fleming [20], and has been investigated in many different situations since then [13, 15, 18, 22, 36]. In this approach, the quantity of interest [e.g., $-(\log P\{S_n/n \in A\})/n$ in Cramér’s Theorem] is represented as the minimal cost U^n for a stochastic control problem. The large deviation rate also has a representation as the minimal cost U for a control problem, though in the case of Cramér’s Theorem it is a deterministic control problem. The large deviation asymptotics then correspond to the convergence of minimal costs: $U^n \rightarrow U$.

Because the limit control problem that defines U is deterministic (at least in the case of Cramér’s Theorem), the use of open loop controls as “comparison controls” is acceptable for the purpose of proving $U^n \rightarrow U$. It is for this reason that simple changes of measure can be used to prove the large deviation lower bound in the traditional change of measure argument. In other words, one can always find a “nearly optimal” control from among the class of open loop controls as $n \rightarrow \infty$. However, in analyzing the optimality of importance sampling schemes we must study the asymptotics of a small noise *stochastic game*, rather than a control problem. The connection with the game is introduced in the next section, and further discussed in Sections 3.1 and 3.2. In the setting of stochastic games, and even for small noise stochastic games, open loop controls are not “nearly optimal” except in special circumstances. As a consequence, to come close to optimality in a general situation one must consider importance sampling schemes that adapt the sampling distribution in the course of simulating each trajectory indexed by k .

The paper is organized as follows. In Section 2 we give the definition of asymptotic optimality, and show that adaptive importance sampling schemes designed to minimize the second moment are asymptotically optimal. Section 3 discusses an alternative formal PDE approach to the adaptive scheme, and describes a method for the construction of a single asymptotically optimal adaptive scheme. Two numerical examples are also included in Section 3.3. Certain technical proofs are deferred to the Appendix to ease exposition.

2 Statement and Proof of the Main Result

Consider a probability space (Ω, \mathcal{F}, P) and a family of events $\{A_n\}$ with

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log P\{A_n\} = -\gamma,$$

for some $\gamma \geq 0$. A general formulation of importance sampling for this problem can be described as follows. In order to estimate $P\{A_n\}$, a generic random variable \bar{Z}_n is constructed such that $P\{A_n\} = E\bar{Z}_n$. Independent replications $(\bar{Z}_n^0, \bar{Z}_n^1, \dots, \bar{Z}_n^{K-1})$ of \bar{Z}_n are then generated, and we obtain an estimator by averaging:

$$\bar{Q}_n^K \doteq \frac{\bar{Z}_n^0 + \bar{Z}_n^1 + \dots + \bar{Z}_n^{K-1}}{K}.$$

The estimator is unbiased, i.e., $E\bar{Q}_n^K = P\{A_n\}$. The rate of convergence associated with this estimator is determined by the variance of the summands, or equivalently, their second moment $E[(\bar{Z}_n^0)^2]$. The smaller the second moment, the faster the convergence, whence the smaller sample size K required. However, it follows from Jensen's inequality that

$$\limsup_{n \rightarrow \infty} -\frac{1}{n} \log E[(\bar{Z}_n^0)^2] \leq \lim_{n \rightarrow \infty} -\frac{1}{n} \log (E\bar{Z}_n^0)^2 = 2\gamma.$$

The estimator \bar{Q}_n^K is said to be *asymptotically optimal* if the upper bound is achieved, i.e., if

$$\lim_{n \rightarrow \infty} -\frac{1}{n} \log E[(\bar{Z}_n^0)^2] = 2\gamma. \quad (2.1)$$

In order to illustrate the main idea we will return, for the remainder of the paper, to the setup of Cramér's theorem.

Remark 2.1 Since the performance of the estimator \bar{Q}_n^K is completely determined by the second moment of its generic, iid building block \bar{Z}_n^k , we will drop the superscript k hereafter. Note that n does *not* stand for sample size. Also, unless noted otherwise, the integral sign will always denote an integral over \mathbb{R}^d .

Let a probability measure μ on \mathbb{R}^d be given, and define H and L by equations (1.1) and (1.2), respectively. Let $A \subset \mathbb{R}^d$ be given. We wish to estimate the probability $p_n \doteq P\{S_n/n \in A\}$, and make use of the following assumption.

Condition 2.1 $H(\alpha) < \infty$ for all $\alpha \in \mathbb{R}^d$, and equation (1.3) holds.

Under this assumption,

$$\lim_{n \rightarrow \infty} -\frac{1}{n} \log p_n = \inf_{\beta \in A^o} L(\beta) = \inf_{\beta \in \bar{A}} L(\beta) = \inf_{\beta \in A} L(\beta). \quad (2.2)$$

We consider a family of changes of measure which are related to μ by an exponential tilt. To this end, we introduce the following control problem. Let a collection of Borel measurable functions (controls) $\alpha^n \doteq \{\alpha_j^n(\cdot), j = 0, 1, \dots, n-1\}$ be given. Then the state dynamics are governed by

$$\bar{S}_j^n = \sum_{i=0}^{j-1} \bar{Y}_i^n, \quad j = 0, 1, \dots, n-1,$$

where \bar{Y}_j^n is conditionally distributed, given $\{\bar{Y}_i^n, i = 0, 1, \dots, j-1\}$, according to

$$v_j^n(dy) = \exp \left\{ \langle \alpha_j^n(\bar{S}_j^n/n), y \rangle - H(\alpha_j^n(\bar{S}_j^n/n)) \right\} \mu(dy).$$

In other words, the sampling distribution is allowed to depend on the historical empirical mean. The estimator of $p_n = P\{S_n/n \in A\}$ is defined as the average of independent replications of

$$\bar{X}_n \doteq 1_{\{\bar{S}_n^n/n \in A\}} e^{\sum_{j=0}^{n-1} (-\langle \alpha_j^n(\bar{S}_j^n/n), \bar{Y}_j^n \rangle + H(\alpha_j^n(\bar{S}_j^n/n)))}.$$

Our goal is to minimize the second moment, hence the variance, of the summands \bar{X}_n by judiciously choosing the control α^n . Thus we consider the value function defined by

$$V^n \doteq \inf_{\alpha^n} E[\bar{X}_n^2] = \inf_{\alpha^n} E \left[1_{\{\bar{S}_n^n/n \in A\}} e^{\sum_{j=0}^{n-1} (-2\langle \alpha_j^n(\bar{S}_j^n/n), \bar{Y}_j^n \rangle + 2H(\alpha_j^n(\bar{S}_j^n/n)))} \right].$$

We also consider the log transform

$$W^n = -\frac{1}{n} \log V^n. \quad (2.3)$$

The following result asserts the asymptotic optimality (or near asymptotic optimality) of the dynamic change of measure defined by any minimizing sequence of controls (respectively, nearly minimizing sequence of controls) in the definition of V^n . Of course, one would like to define a change of measure that does not explicitly depend on n . Such a change of measure is naturally suggested by the deterministic control problem that characterizes the limit of the V^n , and we will remark further on this in Section 3.2.

Theorem 2.1 *Assume Condition 2.1, and define W^n by (2.3). Then*

$$\lim_{n \rightarrow \infty} W^n = 2 \inf_{\beta \in A} L(\beta).$$

The proof of the theorem calls for the following min/max lemma, which is of independent interest. Recall that for a probability distribution γ on \mathbb{R}^d , the *relative entropy* $R(\gamma\|\mu)$ is defined by

$$R(\gamma\|\mu) \doteq \int_{\mathbb{R}^d} \log \frac{d\gamma}{d\mu} d\gamma$$

if $\gamma \ll \mu$, and $R(\gamma\|\mu) \doteq \infty$ if otherwise. Let

$$\mathcal{C} \doteq \left\{ \gamma : \gamma \text{ is a probability distribution on } \mathbb{R}^d, \text{ and } R(\gamma\|\mu) < \infty \right\}.$$

The convexity of relative entropy in the first argument [13, Lemma 1.4.3] implies that \mathcal{C} is a convex set.

Lemma 2.2 *Suppose that $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is a bounded measurable function. Then*

$$\begin{aligned} & \sup_{\alpha \in \mathbb{R}^d} \inf_{\gamma \in \mathcal{C}} \left[\int f(y) \gamma(dy) + R(\gamma\|\mu) + \int \langle \alpha, y \rangle \gamma(dy) - H(\alpha) \right] \\ &= \inf_{\gamma \in \mathcal{C}} \sup_{\alpha \in \mathbb{R}^d} \left[\int f(y) \gamma(dy) + R(\gamma\|\mu) + \int \langle \alpha, y \rangle \gamma(dy) - H(\alpha) \right]. \end{aligned}$$

The proof of this lemma is deferred to the Appendix.

Proof of Theorem 2.1. The unbiasedness of the estimator associated with \bar{X}_n and Jensen's inequality imply that

$$V^n \geq \inf_{\alpha^n} (E \bar{X}_n)^2 = \inf_{\alpha^n} p_n^2 = p_n^2.$$

Equation (2.2) and the last display then give

$$\limsup_{n \rightarrow \infty} W^n \leq \lim_{n \rightarrow \infty} -\frac{1}{n} \log p_n^2 = 2 \inf_{\beta \in A} L(\beta).$$

It remains to show the reverse inequality

$$\liminf_{n \rightarrow \infty} W^n \geq 2 \inf_{\beta \in A} L(\beta).$$

We will use the weak convergence approach as developed in [13]. To analyze the asymptotics of W^n , we represent it as the value function for a stochastic control problem. To this end, we first extend the state dynamics so as to write down a dynamic programming equation (DPE). The “state variable” will be the normalized quantity \bar{S}_i^n/n . Abusing notation a bit, for $x \in \mathbb{R}^d$, $i \in \{0, \dots, n\}$, define the state dynamics

$$\bar{S}_{i,j}^n = nx + \sum_{\ell=i}^{j-1} \bar{Y}_{i,\ell}^n, \quad j = i, \dots, n.$$

Here $\bar{Y}_{i,j}^n$ is conditionally distributed, given $\{\bar{Y}_{i,\ell}^n, \ell = i, \dots, j-1\}$, according to

$$v_{i,j}^n(dy) = \exp \left\{ \left\langle \alpha_j^n(\bar{S}_{i,j}^n/n), y \right\rangle - H \left(\alpha_j^n(\bar{S}_{i,j}^n/n) \right) \right\} \mu(dy).$$

Similarly, define

$$V^n(x, i) \doteq \inf_{\alpha^n} E \left[1_{\{\bar{S}_{i,n}^n/n \in A\}} e^{\sum_{j=i}^{n-1} (-2 \langle \alpha_j^n(\bar{S}_{i,j}^n/n), \bar{Y}_{i,j}^n \rangle + 2H(\alpha_j^n(\bar{S}_{i,j}^n/n)))} \right] \quad (2.4)$$

and

$$W^n(x, i) = -\frac{1}{n} \log V^n(x, i). \quad (2.5)$$

Note that the original case corresponds to $x = 0$ and $i = 0$, that is

$$V^n = V^n(0, 0), \quad W^n = W^n(0, 0), \quad \bar{S}_n^n = \bar{S}_{0,n}^n.$$

Also observe the terminal conditions

$$V^n(x, n) = 1_A(x), \quad W^n(x, n) = \infty \cdot 1_A(x).$$

Since it is inconvenient to study the problem with an ∞ terminal condition, we instead work with a mollified version of the control problem. Let $F : \mathbb{R}^d \rightarrow \mathbb{R}$ be an arbitrary bounded and continuous function. Suppose that V_F^n is defined as in (2.4), save that the indicator function $1_{\{\bar{S}_{i,n}^n/n \in A\}}$ is replaced by $\exp\{-2nF(\bar{S}_{i,n}^n/n)\}$. Similarly define

$$W_F^n(x, i) \doteq -\frac{1}{n} \log V_F^n(x, i).$$

It is not difficult to see that

$$|W_F^n(x, i)| \leq 2\|F\|_\infty, \quad e^{-2n\|F\|_\infty} \leq V_F^n(x, i) \leq e^{2n\|F\|_\infty}. \quad (2.6)$$

Indeed, the first inequality is implied by the second inequality, while the latter is implied by

$$\inf_{\alpha^n} E \left[e^{\sum_{j=i}^{n-1} (-2\langle \alpha_j^n(\bar{S}_{i,j}^n/n), \bar{Y}_{i,j}^n \rangle + 2H(\alpha_{i,j}^n(\bar{S}_{i,j}^n/n)))} \right] = 1.$$

To see that this is true, first consider the case $\alpha^n \equiv 0$. Since $H(0) = 0$, the right hand side is bounded below by the left hand side. Next consider any control sequence α^n . By Jensen's inequality

$$\begin{aligned} & E \left[e^{\sum_{j=i}^{n-1} (-2\langle \alpha_j^n(\bar{S}_{i,j}^n/n), \bar{Y}_{i,j}^n \rangle + 2H(\alpha_{i,j}^n(\bar{S}_{i,j}^n/n)))} \right] \\ & \geq E \left[e^{\sum_{j=i}^{n-1} (-\langle \alpha_j^n(\bar{S}_{i,j}^n/n), \bar{Y}_{i,j}^n \rangle + H(\alpha_{i,j}^n(\bar{S}_{i,j}^n/n)))} \right]^2 \\ & = 1. \end{aligned}$$

Thus the left hand side is bounded below by the right hand side.

It follows from [5] that V_F^n satisfies the Bellman equation

$$\begin{aligned} V_F^n(x, i) &= \inf_{\alpha} \int e^{-2\langle \alpha, y \rangle + 2H(\alpha)} V_F^n \left(x + \frac{1}{n}y, i+1 \right) e^{\langle \alpha, y \rangle - H(\alpha)} \mu(dy) \\ &= \inf_{\alpha} \int e^{-\langle \alpha, y \rangle + H(\alpha)} V_F^n \left(x + \frac{1}{n}y, i+1 \right) \mu(dy), \end{aligned}$$

together with the terminal condition $V_F^n(x, n) = \exp\{-2nF(x)\}$. It follows that

$$W_F^n(x, i) = -\frac{1}{n} \log \inf_{\alpha} \int_{\mathbb{R}^d} e^{-\langle \alpha, y \rangle + H(\alpha)} e^{-nW_F^n(x + \frac{1}{n}y, i+1)} \mu(dy),$$

and that $W_F^n(x, n) = 2F(x)$. The functions V_F^n and W_F^n are both Borel-measurable (see the Appendix for the proof).

The relative entropy representation for exponential integrals [13, Lemma 1.4.2] states that

$$-\log \int e^{-f(x)} \mu(dx) = \inf_{\gamma \in \mathcal{C}} \left[R(\gamma \| \mu) + \int f d\gamma \right]$$

for all bounded measurable functions f . A direct application of this representation to W_F^n is not valid since the function $y \rightarrow \langle \alpha, y \rangle$ is not bounded. Nonetheless, an extension to cover the types of unbounded functions we must consider is possible, and a proof of this fact is given as Lemma 4.1 in

the Appendix. Applying this extended version of the representation gives

$$W_F^n(x, i) = \sup_{\alpha \in \mathbb{R}^d} \inf_{\gamma \in \mathcal{C}} \left[\int W_F^n \left(x + \frac{1}{n}y, i + 1 \right) \gamma(dy) + \frac{1}{n} \left(R(\gamma \parallel \mu) + \int \langle \alpha, y \rangle \gamma(dy) - H(\alpha) \right) \right]. \quad (2.7)$$

The last display shows that W_F^n has an interpretation as the value function for a small noise stochastic game. We will return to this point in Subsection 3.1. The functional on the right-hand side is concave with respect to α and convex with respect to γ , and the sets \mathbb{R}^d and \mathcal{C} are both convex. Since the set \mathcal{C} is in general non-compact, the min/max theorem [38] cannot be applied directly. However, Lemma 2.2 shows that the interchange of inf and sup is still valid in this case. In other words,

$$\begin{aligned} W_F^n(x, i) &= \inf_{\gamma \in \mathcal{C}} \sup_{\alpha \in \mathbb{R}^d} \left[\int W_F^n \left(x + \frac{1}{n}y, i + 1 \right) \gamma(dy) + \frac{1}{n} \left(R(\gamma \parallel \mu) + \int \langle \alpha, y \rangle \gamma(dy) - H(\alpha) \right) \right] \\ &= \inf_{\gamma \in \mathcal{C}} \left[\int W_F^n \left(x + \frac{1}{n}y, i + 1 \right) \gamma(dy) + \frac{1}{n} \left(R(\gamma \parallel \mu) + L \left(\int y \gamma(dy) \right) \right) \right]. \end{aligned} \quad (2.8)$$

Equation (2.8) implies that $W_F^n(x, i)$ also has an interpretation as the minimal cost of a stochastic control problem of the same general form as before. To simplify the notation, we state the control problem only for the case $i = 0$. The control problem will be defined on a probability space $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{P})$, and \tilde{E}_x will denote that the initial condition of the state process is x . An admissible control is a sequence $\{\gamma_j^n, j = 0, 1, \dots, n-1\}$, with each γ_j^n be a stochastic kernel on \mathbb{R}^d given \mathbb{R}^d . Given an admissible control sequence, the state dynamics are defined by $\tilde{S}_0^n = nx$ and

$$\tilde{S}_{j+1}^n \doteq \tilde{S}_j^n + \tilde{Y}_j^n,$$

where

$$\tilde{P} \left\{ \tilde{Y}_j^n \in dy \mid \tilde{Y}_i^n, 0 \leq i < j \right\} = \tilde{P} \left\{ \tilde{Y}_j^n \in dy \mid \tilde{S}_j^n/n \right\} = \gamma_j^n(dy \mid \tilde{S}_j^n/n).$$

We then define value function

$$v_F^n(x, 0) \doteq \inf_{\{\gamma_j^n\}} \tilde{E}_x \left[\sum_{j=0}^{n-1} \frac{1}{n} \left(R(\gamma_j^n \parallel \mu) + L \left(\int y \gamma_j^n(dy) \right) \right) + 2F(\tilde{S}_n^n/n) \right],$$

where the infimum is over all controls $\{\gamma_j^n\}$ and resulting controlled processes $\{\tilde{S}_j^n/n\}$ that start at x at time 0. Since v_F^n also satisfies the DPE (2.8) [5, Chapter 8] and terminal condition $v_F^n(x, n) = W_F^n(x, n) = 2F(x)$, we obtain by induction that $W_F^n(x, i) = v_F^n(x, i)$ for all $x \in \mathbb{R}^d$ and $i \in \{0, \dots, n\}$.

Define a stochastic kernel γ^n on \mathbb{R}^d given $[0, 1]$ by

$$\gamma^n(dy|t) \doteq \begin{cases} \gamma_j^n(dy) & \text{if } t \in [j/n, (j+1)/n), j = 0, 1, \dots, n-2 \\ \gamma_{n-1}^n(dy) & \text{if } t \in [(n-1)/n, 1] \end{cases},$$

and (abusing notation a bit) define the process $\tilde{S}^n = \{\tilde{S}^n(t), t \in [0, 1]\}$ as the piecewise linear interpolation of the controlled sequence $\{\tilde{S}_j^n\}$. Let λ denote Lebesgue measure. Then the definition of $\gamma^n(dy|t)$ and the convexity of L imply that

$$\begin{aligned} & W_F^n(x, 0) \\ &= v_F^n(x, 0) \\ &= \inf_{\{\gamma_j^n\}} \tilde{E}_x \left[\int_0^1 R(\gamma^n(\cdot|t)|\mu) dt + \sum_{j=0}^{n-1} \frac{1}{n} L \left(\int y \gamma_j^n(dy) \right) + 2F(\tilde{S}_n^n/n) \right] \\ &\geq \inf_{\{\gamma_j^n\}} \tilde{E}_x \left[\int_0^1 R(\gamma^n(\cdot|t)|\mu) dt + L \left(\sum_{j=0}^{n-1} \frac{1}{n} \int y \gamma_j^n(dy) \right) + 2F(\tilde{S}_n^n/n) \right] \\ &= \inf_{\{\gamma_j^n\}} \tilde{E}_x \left[R(\gamma^n|\mu \otimes \lambda) + L \left(\int_0^1 \int_{\mathbb{R}^d} y \gamma^n(dy \times dt) \right) + 2F(\tilde{S}_n^n/n) \right], \end{aligned}$$

where $\gamma^n(dy \times dt) \doteq \gamma^n(dy|t)dt$ and $(\mu \otimes \lambda)(dy \times dt) \doteq \mu(dy)dt$. A straightforward weak convergence approach will be adopted to derive the desired inequality (2.9) below. Since the proof is essentially the same as [13, Theorem 5.3.5], we only give a sketch.

For each $\varepsilon > 0$, there exist a sequence of controls $\{\gamma^n, n \in \mathbb{N}\}$ such that

$$W_F^n(x, 0) + \varepsilon \geq \tilde{E}_x \left[R(\gamma^n|\mu \otimes \lambda) + L \left(\int_0^1 \int_{\mathbb{R}^d} y \gamma^n(dy \times dt) \right) + 2F(\tilde{S}_n^n/n) \right]$$

for every n . Furthermore, since L is non-negative and F is bounded, $\{\gamma^n\}$ is indeed uniformly integrable in the sense of [13, Proposition 5.3.2]. For any subsequence of $\{\gamma^n, n \in \mathbb{N}\}$, we can extract a weakly convergent sub-subsequence, still denoted by $\{\gamma^n\}$, such that $\gamma^n \rightharpoonup \gamma$ for some stochastic kernel γ whose second marginal is Lebesgue measure. We utilize the Skorokhod representation [6], which allows us to assume (when calculating the limits of the integrals) that the convergence is actually w.p.1. It follows

from the lower semicontinuity of R , the convergence $\gamma^n \Rightarrow \gamma$, and the uniform integrability that w.p.1

$$\liminf_n R(\gamma^n \| \mu \otimes \lambda) \geq R(\gamma \| \mu \otimes \lambda),$$

$$\lim_n \int_0^1 \int_{\mathbb{R}^d} y \gamma^n(dy \times dt) = \int_0^1 \int_{\mathbb{R}^d} y \gamma(dy \times dt),$$

and

$$\tilde{S}^n/n - Z \rightarrow 0, \quad \text{where } Z(t) \doteq x + \int_0^t \int_{\mathbb{R}^d} y \gamma(dy|s) ds,$$

where $\gamma(dy \times ds) = \gamma(dy|s) ds$. The convergence in the last display is with respect to the supremum norm. In addition, the convexity of $R(\nu \| \mu)$ in ν and Jensen's inequality imply

$$R(\gamma \| \mu \otimes \lambda) = \int_0^1 R(\gamma(\cdot|t) \| \mu(\cdot)) dt \geq R\left(\int_0^1 \gamma(\cdot|t) dt \| \mu(\cdot)\right).$$

The uniform integrability of [13, Theorem 5.3.5] also implies $\int_0^1 \int_{\mathbb{R}^d} \|y\| \gamma(dy \times dt) < \infty$ w.p.1, and hence

$$\int_0^1 \int_{\mathbb{R}^d} y \gamma(dy|t) dt = \int_{\mathbb{R}^d} \int_0^1 y \gamma(dy|t) dt.$$

By Fatou's Lemma and the lower-semicontinuity of L [30], we have

$$\begin{aligned} & \liminf_n W_F^n(x, 0) + \varepsilon \\ & \geq \tilde{E}_x \left[R\left(\int_0^1 \gamma(\cdot|t) dt \| \mu(\cdot)\right) + L\left(\int_{\mathbb{R}^d} \int_0^1 y \gamma(dy|t) dt\right) + 2F(Z(1)) \right] \end{aligned}$$

Using the identity

$$\inf \left[R(\nu \| \mu) : \int y \nu(dy) = \beta \right] = L(\beta)$$

[13, Lemma 3.3.3], it is elementary to show that the right hand side of the last inequality is bounded below by

$$2 \inf_{\beta \in \mathbb{R}^d} [L(\beta) + F(x + \beta)].$$

Since $\varepsilon > 0$ is arbitrary,

$$\liminf_{n \rightarrow \infty} W_F^n(x, 0) \geq 2 \inf_{\beta \in \mathbb{R}^d} [L(\beta) + F(x + \beta)],$$

and in particular

$$\liminf_{n \rightarrow \infty} W_F^n(0, 0) \geq 2 \inf_{\beta \in \mathbb{R}^d} [L(\beta) + F(\beta)]. \quad (2.9)$$

Now let $F_j(y) \doteq j(d(y, \bar{A}) \wedge 1)$. Since $1_A(y) \leq \exp\{-2nF_j(y)\}$,

$$\begin{aligned} \liminf_{n \rightarrow \infty} W^n &= \liminf_{n \rightarrow \infty} W^n(0, 0) \\ &\geq \liminf_{n \rightarrow \infty} W_{F_j}^n(0, 0) \\ &\geq 2 \inf_{\beta \in \mathbb{R}^d} [L(\beta) + F_j(\beta)]. \end{aligned}$$

Exactly as on [13, pages 10-11], a compactness argument shows that

$$\lim_{j \rightarrow \infty} \inf_{\beta \in \mathbb{R}^d} \{L(\beta) + F_j(\beta)\} = \inf_{\beta \in \bar{A}} L(\beta),$$

and we complete the proof. ■

For a scalar a let $\lfloor a \rfloor$ denote the integer part of a .

Proposition 2.1 Assume Condition 2.1, and define $W^n(x, i)$ as in (2.5), except in (1.3) we replace A with $(A - x)/(1 - t)$. Then

$$\lim_{n \rightarrow \infty} W^n(x, \lfloor tn \rfloor) \rightarrow 2U(x, t),$$

where

$$U(x, t) \doteq \inf\{(1 - t)L(\beta) : x + (1 - t)\beta \in A\}.$$

Proof. The proof is essentially the same and thus omitted. ■

Remark 2.2 The proof of Theorem 2.1 actually implies a more general result. If we write $W_F^n \doteq W_F^n(0, 0)$, then

$$\lim_{n \rightarrow \infty} W_F^n = 2 \inf_{\beta \in \mathbb{R}^d} [L(\beta) + F(\beta)]$$

for all bounded and continuous functions $F : \mathbb{R}^d \rightarrow \mathbb{R}$. Indeed, since (2.9) gives the reverse inequality, all we need to show is that

$$\limsup_n W_F^n \leq 2 \inf_{\beta \in \mathbb{R}^d} [L(\beta) + F(\beta)].$$

However, by definition

$$W_F^n \doteq -\frac{1}{n} \log V_F^n,$$

where $V_F^n \doteq V_F^n(0, 0)$ is defined by

$$V_F^n = \inf_{\alpha^n} E \left[(\bar{X}_{n,F})^2 \right]$$

with

$$\bar{X}_{n,F} \doteq e^{-nF(\bar{S}_n/n)} e^{\sum_{j=0}^{n-1} (-\langle \alpha_j^n(\bar{S}_j^n/n), \bar{Y}_j^n \rangle + H(\alpha_j^n(\bar{S}_j^n/n)))}.$$

For each control α^n , it is not difficult to verify that the resulting $\bar{X}_{n,F}$ is an unbiased estimator for $E[\exp\{-nF(S_n/n)\}]$. It then follows from Jensen's inequality that

$$W_F^n \leq -\frac{1}{n} \log \inf_{\alpha^n} E \left[(\bar{X}_{n,F})^2 \right] = -\frac{2}{n} \log E e^{-nF(S_n/n)}.$$

Thanks to the large deviations properties of S_n/n , this implies [13, Theorem 1.2.1]

$$\limsup_n W_F^n \leq \limsup_n -\frac{2}{n} \log E e^{-nF(S_n/n)} \leq 2 \inf_{\beta \in \mathbb{R}^d} [L(\beta) + F(\beta)].$$

Similar to Proposition 2.1, we have the following more general version.

$$\lim_n W_F^n(x, \lfloor tn \rfloor) = 2U_F(x, t)$$

where

$$U_F(x, t) \doteq \inf_{\beta \in \mathbb{R}^d} \{(1-t)L(\beta) + F(x + (1-t)\beta)\}. \quad (2.10)$$

3 Examples and Further Remarks

3.1 The formal PDE approach and the limit control problem

Although Theorem 2.1 establishes that there are asymptotically optimal [in the sense of equation (2.1)] adaptive importance sampling schemes, it does not explicitly discuss the construction of such schemes, or even schemes that are approximately asymptotically optimal.

In this subsection, we will discuss an alternative PDE approach to study the adaptive importance sampling scheme. The approach, though largely formal, will shed light on how to construct asymptotically optimal and nearly optimal controls α^n . Some numerical confirmation of this approach is presented in Section 3.3.

For a general measurable function F , and assuming the limit exists, define

$$W_F(x, t) \doteq \lim_n W_F^n(x, \lfloor tn \rfloor).$$

Remark 2.2 asserts that

$$W_F(x, t) = 2U_F(x, t) \quad (3.1)$$

for all continuous and bounded functions F , while Proposition 2.1 confirms the validity of (3.1) for $F = \infty 1_A$. We will re-derive equation (3.1) below by formally arguing that W_F and $2U_F$ satisfy the same PDE.

To obtain this PDE, we rewrite (2.7) as

$$0 = \sup_{\alpha \in \mathbb{R}^d} \inf_{\gamma \in \mathcal{C}} \left[\int \Delta W_F^n(y) \gamma(dy) + \frac{1}{n} \left(R(\gamma \| \mu) + \int \langle \alpha, y \rangle \gamma(dy) - H(\alpha) \right) \right],$$

where

$$\Delta W_F^n(y) \doteq W_F^n \left(x + \frac{1}{n} y, i + 1 \right) - W_F^n(x, i).$$

Suppose that a t subscript denotes the partial derivative with respect to t , and that an x subscript denotes the vector of partials with respect to $x_i, i = 1, \dots, d$. Formally, we have the approximation

$$\Delta W_F^n \approx \frac{1}{n} (W_F)_t + \frac{1}{n} \langle (W_F)_x, y \rangle.$$

Inserting this into the DPE leads to

$$0 = \sup_{\alpha \in \mathbb{R}^d} \inf_{\gamma \in \mathcal{C}} \left[(W_F)_t + \int \langle (W_F)_x, y \rangle \gamma(dy) + R(\gamma \| \mu) + \int \langle \alpha, y \rangle \gamma(dy) - H(\alpha) \right].$$

We have already noted that

$$\inf \left[R(\gamma \| \mu) : \int y \gamma(dy) = \beta \right] = L(\beta).$$

Therefore,

$$0 = (W_F)_t + \sup_{\alpha \in \mathbb{R}^d} \inf_{\beta \in \mathbb{R}^d} [\langle (W_F)_x, \beta \rangle + L(\beta) + \langle \alpha, \beta \rangle - H(\alpha)].$$

Such an equation is called an *Isaac's equation*, and it is well known that such equations are satisfied by the value function for a differential game. We will not give the rather lengthy and detailed formal definition of the game here, but simply note its general features. Because of the intervening minus sign, the maximizing player is actually trying to minimize the variance through the choice of control $\alpha(t) \in \mathbb{R}^d$. The minimizing player appears due to the large deviations approximation of the variance, and chooses $\beta(t) \in$

\mathbb{R}^d . The dynamics $\dot{\phi}(t) = \beta(t)$ only involve the minimizing player, the running cost $L(\beta) + \langle \alpha, \beta \rangle - H(\alpha)$ involves both, and the terminal cost is $2F(\phi(1))$. The adaptive importance sampling scheme will be defined in terms of the optimal control for the maximizing player, which one would prefer to obtain in feedback form. Representing the infimum in terms of the Legendre transform H of L gives

$$0 = (W_F)_t + \sup_{\alpha \in \mathbb{R}^d} [-H(-\alpha - (W_F)_x) - H(\alpha)].$$

The strict convexity of H implies that the supremum is uniquely achieved at

$$\alpha^*(x, t) = -(W_F)_x(x, t)/2, \quad (3.2)$$

and that W_F should satisfy

$$0 = (W_F)_t - 2H(-(W_F)_x/2). \quad (3.3)$$

Lastly, we observe that W_F should also formally satisfy with terminal condition $W_F(x, 1) = 2F(x)$.

Next we formally derive the PDE for U_F . Owing to the convexity of L , U_F is the value function of the deterministic control problem

$$U_F(x, t) = \inf_{\phi} \left[\int_t^1 L(\dot{\phi}(s)) ds + F(\phi(1)) \right], \quad (3.4)$$

where the infimum is over all absolutely continuous ϕ which satisfy $\phi(t) = x$. A standard dynamic programming argument implies that U_F satisfies [4, 21] the DPE

$$0 = (U_F)_t + \inf_{a \in \mathbb{R}^d} [L(a) + \langle a, (U_F)_x \rangle] = (U_F)_t - H(-(U_F)_x), \quad (3.5)$$

with terminal condition $U_F(x, 1) = F(x)$.

Comparing the PDE (3.3) with (3.5), and noting $W_F(x, 1) = 2U_F(x, 1)$, we conclude that (3.1) holds if either PDE has a unique solution. Furthermore, note that α^* defined by (3.2) also satisfies

$$\alpha^*(x, t) = -(W_F)_x(x, t)/2 = -(U_F)_x(x, t). \quad (3.6)$$

3.2 Implementation issues

One approach to the construction of optimal or nearly optimal adaptive importance sampling schemes (i.e., selection of the control α^n), would be to

solve (numerically if need be) the DPE associated with W_F^n . However, any such scheme would directly depend on n , and in general, one would prefer schemes without this dependence.

An alternative approach, which is to be discussed in this subsection, is to consider the DPE associated with the limit problem W_F (equivalently U_F). In the preceding subsection, we formally characterized W_F as the solution to an Isaac's equation, and hence as the value function for a differential game. We also showed that W_F can be characterized as the value function of a deterministic optimal control problem, which is often easier to solve or approximate numerically.

The equation (3.6) identifies, at least formally, an optimal feedback control policy. However, this observation is not totally satisfactory for several reasons. The first is that even if we have an exact formula for U_F (or W_F), the DPE (3.5) is usually satisfied only in a weak sense, in which case the partial derivative $(U_F)_x$ may not be defined for all times and spatial points. Under additional conditions one can show that the set of points where $\alpha^*(x, t)$ is well defined is open and dense [19]. At points where the gradient is not defined there are often superdifferentials, and a natural replacement for the DPE suggests feedback controls defined in terms of extreme superdifferentials rather than gradients. Nonetheless, a comprehensive theory is not available. The second reason is that U_F (or, W_F) does not usually have an explicit solution, and so in many cases a numerical approximation is required. Convergent numerical approximations have been studied extensively (see, e.g., [27]). Practical experience and some theory (e.g., [16]), have shown that the numerical schemes which construct (provably) convergent approximations to U_F (or, W_F) also yield nearly optimal feedback controls, though they are usually limited by practical implementation constraints to low dimensional problems.

Our goal in this subsection is to formally characterize the optimal control α^* at all points (x, t) through the dual relation in the Legendre transform (1.2). It is straightforward to see from the control problem (3.4) that an optimal control at (x, t) is the minimizer in equation (2.10), say $\beta^*(x, t)$, thanks to the convexity of L . Note that the existence of $\beta^*(x, t)$ is guaranteed with very mild conditions, e.g., F is bounded and continuous. This implies that $-(U_F)_x(x, t)$ and $\beta^*(x, t)$ are conjugate. It follows from (3.6) that

$$\alpha^*(x, t) \text{ is conjugate to the minimizer } \beta^*(x, t) \text{ in (2.10).}$$

At points where $(U_F)_x(x, t)$ exists this definition gives $\alpha^*(x, t) = -(U_F)_x(x, t)$. At points where $(U_F)_x(x, t)$ does not exist there are multiple minimizing

$\beta^*(x, t)$, and one should define $\alpha^*(x, t)$ though conjugacy in any Borel measurable way.

Remark 3.1 In the setting of Cramér’s theorem, $F = \infty \cdot 1_A$. In this case,

$$\beta^*(x, t) \in \operatorname{argmin}\{(1-t)L(\beta) : x + (1-t)\beta \in A\},$$

and $\alpha^*(x, t)$ is its conjugate. This also implies that, at time (x, t) , the change of measure associated with $\alpha^*(x, t)$ actually shifts the mean to $\beta^*(x, t)$. Roughly speaking, this means that the controlled random walk always points to the “most likely” end point, given that the end point is in A and that we are currently at x with $1-t$ units of time to go. Note that we can assume without loss that A is closed under the conditions of Proposition 2.1. This suffices to guarantee the existence of $\beta^*(x, t)$ for all $x \in \mathbb{R}^d, t \in [0, 1)$.

3.3 The relations between blind and adaptive importance sampling and differential games in the small and the large

In the adaptive sampling scheme, the sampling distribution is allowed to be a function of the current state of the “controlled empirical mean” and the time index. In contrast, the scheme suggested by the standard heuristic in the setting of Cramér’s Theorem is a fixed change of measure throughout. In a setting more general than Cramér’s Theorem (e.g., the models of Chapter 6 of [13]) the scheme suggested by the standard heuristic might also depend on the time index, though not on any controlled state (in other words, it is an “open loop” control). In both cases, a logarithmic transform would be used to characterize the asymptotic behavior of the second moment, and this logarithmic transformation introduces another control through the variational formula for relative entropy (cf. [13] and Section 2).

In the case of adaptive sampling, the control that tries to minimize the second moment maximizes in the logarithmic transform, owing to an intervening minus sign. The control that large deviations introduces attempts to minimize, and an examination of the dynamic programming equation shows that at each time step the “large deviation control” gets to see the variance control used to generate the next sample, thus giving it the “information advantage.”

Thus in the limit $n \rightarrow \infty$ we expect to obtain a differential game with the advantage given to the minimizing player, the so-called “lower game” (see, e.g., [17]). We refer to such games as games “in the small,” since in the

discrete time prelimit game the players' selections of controls are interleaved and sequential.

In the case of the standard heuristic the limit game is quite different. Here the maximizing player (the player who selects the sampling distribution) must show his entire "open loop" control to his opponent, who can then minimize given knowledge of the opponent's control for all $t \in [0, 1]$. Such a game will be referred to as the "lower game in the large." It is easy to show that the "lower game in the large" is never greater than the "lower game in the small." When these two coincide, one would expect that the standard heuristic and the adaptive importance sampling scheme have the same (large deviation) asymptotic properties, even though there could still be a significant difference in performance for any given value of n . However, when there is a gap between the values of the two games the standard heuristic should not be asymptotically optimal, and in fact the relative error in the standard heuristic should grow exponentially with n .

3.4 Examples

In this section we include some numerical examples to illustrate the subtle pitfalls of the blind (i.e., standard heuristic) importance sampling. All the simulations are performed with standard S-Plus software, and the time difference between blind importance sampling and adaptive importance sampling is negligible. Each simulation takes at most a few seconds.

Example 1: Let Y_0, Y_1, \dots be a sequence of iid $N(1, 1)$ random variables, and $S_n = Y_0 + \dots + Y_{n-1}$ the partial sum. Consider the set

$$A = (-\infty, a] \cup [b, \infty)$$

with $a < 1 < b$. We want to estimate $P\{S_n/n \in A\}$ by importance sampling.

It is not difficult to verify that

$$H(\alpha) = \log E \left[e^{\alpha Y_0} \right] = \alpha + \frac{1}{2}\alpha^2, \quad L(\beta) = \sup_{\alpha \in \mathbb{R}} [\alpha\beta - H(\alpha)] = \frac{1}{2}(\beta - 1)^2.$$

It follows from Cramér's theorem that

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log P\{S_n/n \in A\} = - \inf_{\beta \in A} L(\beta).$$

From now on, we will assume $a + b < 2$, which implies that $\beta^* = b$ achieves the minimum of L over A .

The blind importance sampling identifies the change of measure as the exponential tilt associated with the supremizing point α^* in $\alpha \rightarrow \alpha\beta^* - H(\alpha)$, which is

$$\alpha^* = \beta^* - 1 = b - 1.$$

In other words, the new probability measure is

$$\nu(dy) = e^{\alpha^*y - H(\alpha^*)} \mu(dy) = e^{\alpha^*y - H(\alpha^*)} \cdot \frac{1}{\sqrt{2\pi}} e^{-\frac{(y-1)^2}{2}} dy = \frac{1}{\sqrt{2\pi}} e^{-\frac{(y-b)^2}{2}} dy,$$

which is just the probability measure associated with $N(b, 1)$. The algorithm starts by generating iid sequences $\{\bar{Y}_i^k\}$ according to $N(b, 1)$, and then forms the estimator

$$\frac{1}{K} \sum_{k=0}^{K-1} 1_{\{\bar{S}_n^k/n \in A\}} \prod_{i=0}^{n-1} \left[e^{-\alpha^* \bar{Y}_i^k + H(\alpha^*)} \right] = \frac{1}{K} \sum_{k=0}^{K-1} 1_{\{\bar{S}_n^k/n \in A\}} e^{-\alpha^* \bar{S}_n^k + nH(\alpha^*)}.$$

The rate of convergence of this estimator is determined by the second moment of the summands, say \tilde{V}^n , which satisfies (see the discussion in the Introduction)

$$\lim_n -\frac{1}{n} \log \tilde{V}^n = \inf_{\beta \in A} [\alpha^* \beta - H(\alpha^*) + L(\beta)].$$

If β^* achieves the infimum, then the right hand side equals $2L(\beta^*)$, which is the optimal rate for the second moment. However, this is not always true. Indeed, it is not difficult to verify that

$$\lim_n -\frac{1}{n} \log \tilde{V}^n = \frac{1}{2}(a+b-2)^2 - (b-1)^2 < 2L(\beta^*)$$

if $a + 3b > 4$.

In the numerical simulation below, we take $n = 60$, $K = 5000$, $a = 0.75$, $b = 1.2$. The true value of the probability we want to estimate is

$$p_n = 1 - \Phi((b-1)\sqrt{n}) + \Phi((a-1)\sqrt{n}) = 8.7\%.$$

The table below presents the outcome of 4 simulations.

	No. 1	No. 2	No. 3	No. 4
Estimate \hat{p}_n	15.42	7.81	5.96	6.02
Standard Error	4.70	1.75	0.12	0.12
95% Confidence Interval	[6.02, 25.82]	[4.31, 11.31]	[5.72, 6.20]	[5.78, 6.26]
Number of "Rogue" Trajectory	4	1	0	0
$(-\log \tilde{V}^n)/(-\log \hat{p}_n)$	-1.3	-0.14	1.6	1.6

Table 1.

In simulations No. 1 and No. 2, the presence of “rogue” trajectories (see the Introduction) greatly raises the standard errors associated with the estimates. Indeed, in each of these two simulations, the proportion of the estimate of the second moment due to these few “rogue” trajectories is more than 99%. In simulations No. 3 and No. 4, however, there are no “rogue” trajectories, and the standard error associated with the estimate is deceptively small. The reason is that the standard error is itself estimated from the sample. Without “rogue” trajectories, we actually underestimate the standard error. Therefore, we cannot put much confidence in the standard errors thus obtained, or in the “tight” confidence intervals that follow. Note that the confidence intervals from these two simulations do not contain the true value.

The adaptive importance sampling, on the other hand, yields more accurate estimates and its performance is much more stable. Below is the numerical result. The standard errors (almost the same across different simulations) are much smaller compared to those of the blind algorithm.

	No. 1	No. 2	No. 3	No. 4
Estimate \hat{p}_n (%)	8.66	8.59	8.72	8.66
Standard Error (%)	0.06	0.06	0.06	0.07
95% Confidence Interval (%)	[8.54, 8.78]	[8.47, 8.71]	[8.60, 8.84]	[8.52, 8.80]
$(-\log \hat{V}^n)/(-\log \hat{p}_n)$	1.51	1.45	1.49	1.39

Table 2.

The selection of a nearly optimal control $\alpha^n = \{\alpha_j^n(\cdot) : j = 0, 1, \dots, n-1\}$ was discussed in the preceding sections. It was formally argued in Subsection 3.1 that the associated limiting differential game indicates

$$\alpha_j^n(x) = -U_x(x, j/n),$$

is a good choice, where

$$U(x, t) = \inf\{(1-t)L(\beta) : x + (1-t)\beta \in A\}.$$

In this example, it is not difficult to verify that

$$U(x, t) = \begin{cases} 0, & \text{if } x \geq b - (1-t) \text{ or } x \leq a - (1-t) \\ \frac{1-t}{2} \left(\frac{b-x}{1-t} - 1 \right)^2, & \text{if } (a+b)/2 - (1-t) \leq x \leq b - (1-t) \\ \frac{1-t}{2} \left(\frac{a-x}{1-t} - 1 \right)^2, & \text{if } a - (1-t) \leq x \leq (a+b)/2 - (1-t) \end{cases}$$

and

$$-U_x(x, t) = \begin{cases} 0 & , \quad \text{if } x > b - (1 - t) \text{ or } x < a - (1 - t) \\ \frac{b-x}{1-t} - 1, & \text{if } (a+b)/2 - (1-t) < x < b - (1-t) \\ \frac{a-x}{1-t} - 1, & \text{if } a - (1-t) < x < (a+b)/2 - (1-t) \end{cases} .$$

At points where $-U_x$ is not defined any extreme superdifferential may be used in lieu of U_x . In all cases this is the same as using the maximizing α in $\alpha \rightarrow \alpha\beta^*(x, t) - H(\alpha)$, where

$$\beta^*(x, t) \in \operatorname{argmin}\{(1-t)L(\beta) : x + (1-t)\beta \in A\}.$$

See Subsection 3.2 for more discussion.

To summarize, the distribution of \bar{Y}_j^k given $\{\bar{Y}_i^k, i = 0, \dots, j-1\}$ is normal with variance 1 and mean $-U_x(x, j/n) + 1 = \beta^*(x, j/n)$, with $x = (\bar{Y}_0^k + \dots + \bar{Y}_{j-1}^k)/n$.

The following table illustrates the asymptotic optimality of the adaptive importance sampling as $n \rightarrow \infty$.

	$n = 100$	$n = 200$	$n = 500$
Theoretical p_n	2.90×10^{-2}	2.54×10^{-3}	3.88×10^{-6}
Estimate \hat{p}_n	2.98×10^{-2}	2.57×10^{-3}	3.86×10^{-6}
Standard Error	0.06×10^{-2}	0.04×10^{-3}	0.07×10^{-6}
95% Confidence Interval	$[2.86, 3.10] \times 10^{-2}$	$[2.49, 2.65] \times 10^{-3}$	$[3.72, 4.00] \times 10^{-6}$
$(-\log \hat{V}^n)/(-\log \hat{p}_n)$	1.65	1.84	1.93

Table 3.

Example 2: Let $\{Y_i = (Y_i^0, Y_i^1)\}$ be an iid sequence of two-dimensional, normally distributed random vectors with mean $\mathbf{0}$ and covariance matrix \mathbf{I} , and $S_n = Y_0 + \dots + Y_{n-1}$ the partial sum. For $\alpha, \beta \in \mathbb{R}^2$ we have

$$H(\alpha) = \log E \left[e^{\langle \alpha, Y_0 \rangle} \right] = \frac{1}{2} \|\alpha\|^2, \quad L(\beta) = \sup_{\alpha \in \mathbb{R}^2} [\langle \alpha, \beta \rangle - H(\alpha)] = \frac{1}{2} \|\beta\|^2.$$

Consider the set

$$A \doteq \left\{ y = (y^0, y^1) : \|y + (a, 0)\|^2 = (y^0 + a)^2 + (y^1)^2 \geq r^2 \right\},$$

where $0 < a < r$ are constants. Then Cramér's Theorem yields

$$\lim_n \frac{1}{n} \log P\{S_n/n \in A\} = - \inf_{\beta \in A} L(\beta) = -L(\beta^*) = -\frac{1}{2}(r-a)^2,$$

with minimizing $\beta^* = (r - a, 0)$. It is not difficult to identify the supremizing α in $\alpha \rightarrow \langle \alpha, \beta^* \rangle - H(\alpha)$ as

$$\alpha^* = \beta^* = (r - a, 0),$$

and the blind importance sampling will sample under the new probability measure

$$\nu(dy) = e^{\langle \alpha^*, y \rangle - H(\alpha^*)} \mu(dy) = \frac{1}{2\pi} e^{-\frac{\|y - \alpha^*\|^2}{2}} dy,$$

or the distribution of $N(\beta^*, \mathbf{I})$. As before, the estimator is

$$\frac{1}{K} \sum_{k=0}^{K-1} 1_{\{\bar{S}_n^k/n \in A\}} \prod_{i=0}^{n-1} e^{-\langle \alpha^*, \bar{S}_n^k \rangle + nH(\alpha^*)},$$

where $\bar{S}_n^k = \bar{Y}_0^k + \dots + \bar{Y}_{n-1}^k$ and $\{Y_i^k\}$ are iid $N(\beta^*, \mathbf{I})$ random vectors.

It is not difficult to verify that the second moment \tilde{V}^n associated with the summands satisfies

$$\lim_n -\frac{1}{n} \log \tilde{V}^n = \inf_{\beta \in A} [\langle \alpha^*, \beta \rangle - H(\alpha^*) + L(\beta)] = \inf_{\beta \in A} \left[\frac{1}{2} \|\beta + \alpha^*\|^2 - \|\alpha^*\|^2 \right].$$

If we further assume that $0 < a < r/2$, then the infimum is not achieved at β^* , but at $(-r - a, 0)$. In fact, β^* achieves the *maximum* of the right hand side over $\beta \in \partial A$. In this case, we have

$$\lim_n -\frac{1}{n} \log \tilde{V}^n = 2a^2 - (r - a)^2 < 2L(\beta^*) = (r - a)^2.$$

In the table below, we take $n = 60$, $K = 5000$, $r = 0.5$, $a = 0.05$. The theoretical value of the probability can be obtained as follows:

$$p_n \doteq P\{S_n/n \in A\} = P\{\|S_n/\sqrt{n} + (\sqrt{na}, 0)\|^2 \geq nr^2\}.$$

Since S_n/\sqrt{n} is $N(\mathbf{0}, \mathbf{I})$ distributed, $\|S_n/\sqrt{n} + (\sqrt{na}, 0)\|^2$ is a *non-central chi-square* random variable with 2 degree of freedom and noncentrality parameter $\|(\sqrt{na}, 0)\|^2 = na^2$. Its distribution is available in standard statistics softwares like S-Plus. Here we have

$$p_n = 8.97 \times 10^{-4}.$$

	No. 1	No. 2	No. 3	No. 4
Estimate \hat{p}_n ($\times 10^{-4}$)	6.38	6.04	7.26	11.09
Standard Error ($\times 10^{-4}$)	0.78	0.42	1.17	5.40
95% Confidence Interval ($\times 10^{-4}$)	[4.82, 7.94]	[5.20, 6.88]	[4.92, 9.60]	[0.29, 21.89]
$(-\log \tilde{V}^n)/(-\log \hat{p}_n)$	1.41	1.56	1.33	0.96

Table 4.

In this example, we do not have a clear-cut definition of a “rogue” trajectory. The main idea, however, remains the same. There are a continuum of possibilities for exiting away from β^* and building up a large Radon-Nikodym derivative, and to varying degrees these trajectories degrade the estimate. As the table above shows, both the estimated probabilities and estimated second moments vary considerably across different simulations.

The adaptive importance sampling again outperforms the blind algorithm. Below are the simulation results, and the estimates are clearly more accurate and stable.

	No. 1	No. 2	No. 3	No. 4
Estimate $\hat{p}_n (\times 10^{-4})$	9.29	8.85	8.90	9.18
Standard Error ($\times 10^{-4}$)	0.21	0.24	0.18	0.28
95% Confidence Interval ($\times 10^{-4}$)	[8.87, 9.71]	[8.37, 9.33]	[8.54, 9.26]	[8.62, 9.74]
$(-\log \hat{V}^n)/(-\log \hat{p}_n)$	1.82	1.78	1.84	1.75

Table 5.

The choice of control α^n is obtained as before. One can compute the function U explicitly and then let $\alpha_j^n(x) = -U_x(x, j/n)$, as in Example 1. Alternatively, as discussed in Subsection 3.2, we can use that $-U_x$ is the maximizing α in $\langle \alpha, \beta^*(x, t) \rangle - H(\alpha)$, where

$$\beta^*(x, t) = \operatorname{argmin} \{(1-t)L(\beta) : x + (1-t)\beta \in A\}.$$

It is not difficult to check that

$$-U_x(x, t) = \beta^*(x, t) = \begin{cases} \mathbf{0} & , \quad \text{if } x \in A \\ \left(\frac{r}{\|x+(a,0)\|} - 1 \right) \cdot [x + (a, 0)], & \text{if } x \notin A. \end{cases}$$

In other words, the distribution of \bar{Y}_j^k given $\{\bar{Y}_i^k, i = 0, \dots, j-1\}$ is normal with variance \mathbf{I} and mean $\beta^*(x, j/n)$ when $x = (\bar{Y}_0^k + \dots + \bar{Y}_{j-1}^k)/n$.

Below are more simulation results, which illustrate the asymptotic optimality of the adaptive importance sampling as n tends to infinity.

	$n = 40$	$n = 80$	$n = 120$
Theoretical p_n	8.49×10^{-3}	1.00×10^{-4}	1.40×10^{-6}
Estimate \hat{p}_n	8.64×10^{-3}	0.98×10^{-4}	1.39×10^{-6}
Standard Error	0.25×10^{-3}	0.02×10^{-4}	0.03×10^{-6}
95% Confidence Interval	$[8.14, 9.14] \times 10^{-3}$	$[0.94, 1.02] \times 10^{-4}$	$[1.33, 1.45] \times 10^{-6}$
$(-\log \hat{V}^n)/(-\log \hat{p}_n)$	1.65	1.85	1.92

Table 6.

4 Appendix

Lemma 4.1 *Let μ be a probability measure on \mathbb{R}^d and let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be bounded and measurable. Define $H(\alpha)$ by (1.1), and assume $H(\alpha) < \infty$ for all $\alpha \in \mathbb{R}^d$. Then for each $\alpha \in \mathbb{R}^d$,*

$$\begin{aligned} & -\log \int e^{-\langle \alpha, y \rangle + H(\alpha) - f(y)} \mu(dy) \\ &= \inf_{\{\gamma: R(\gamma\|\mu) < \infty\}} \left[R(\gamma\|\mu) + \int f(y) \gamma(dy) + \int \langle \alpha, y \rangle \gamma(dy) - H(\alpha) \right] \end{aligned}$$

Proof. Fix $\alpha \in \mathbb{R}^d$, and define a change of probability measure by

$$\frac{d\theta}{d\mu}(y) \doteq e^{-\langle \alpha, y \rangle - H(-\alpha)}.$$

It follows that

$$\begin{aligned} & -\log \int e^{-\langle \alpha, y \rangle + H(\alpha) - f(y)} \mu(dy) \\ &= -H(\alpha) - H(-\alpha) - \log \int e^{-f(y)} \theta(dy) \\ &= -H(\alpha) - H(-\alpha) + \inf_{\{\gamma: R(\gamma\|\theta) < \infty\}} \left[R(\gamma\|\theta) + \int f(y) \gamma(dy) \right] \end{aligned}$$

Into this expression we insert

$$\begin{aligned} R(\gamma\|\theta) &= \int \log \frac{d\gamma}{d\theta} d\gamma \\ &= \int \log \frac{d\gamma}{d\mu} d\gamma + \int \log \frac{d\mu}{d\theta} d\gamma \\ &= R(\gamma\|\mu) + H(-\alpha) + \int \langle \alpha, y \rangle \gamma(dy). \end{aligned}$$

It remains to show that

$$\{\gamma : R(\gamma\|\theta) < \infty\} = \{\gamma : R(\gamma\|\mu) < \infty\}.$$

Indeed, if $R(\gamma\|\mu) < \infty$ and Condition 2.1 hold then by [13, Lemma 1.4.3] $\int \|y\| \gamma(dy) < \infty$. It then follows that $R(\gamma\|\theta) < \infty$. The proof for the reverse is exactly the same. ■

Proof of the Borel Measurability of V_F^n and W_F^n . We only need to show the Borel-measurability of $V_F^n(x, i)$, and to do this we will use induction

on i . Clearly $V_F^n(x, n) = \exp\{-2nF(x)\}$ is Borel measurable. Suppose that $V_F^n(x, i+1)$ is Borel measurable. Since the infimum of countably many Borel measurable functions is still Borel measurable, it suffices to show that

$$V_F^n(x, i) = \inf_{\alpha \in \mathbb{Q}^d} \int e^{-\langle \alpha, y \rangle + H(\alpha)} V_F^n\left(x + \frac{1}{n}y, i+1\right) \mu(dy),$$

where \mathbb{Q} is the set of rationals. Indeed, for any $\alpha \in \mathbb{R}^d$, there exist a sequence $\{\alpha_m\} \in \mathbb{Q}^d$ such that $\|\alpha_m\| \leq 2\|\alpha\|$ and $\alpha_m \rightarrow \alpha$. Using (2.6) we have the bound

$$e^{-\langle \alpha_m, y \rangle + H(\alpha_m)} V_F^n\left(x + \frac{1}{n}y, i+1\right) \leq e^{2\|\alpha\| \cdot \|y\| + \sup\{H(a) : \|a\| \leq 2\|\alpha\|\}} \cdot e^{2n\|F\|_\infty},$$

and the right hand side is integrable. The Dominated Convergence Theorem implies

$$\int e^{-\langle \alpha_m, y \rangle + H(\alpha_m)} V_F^n\left(x + \frac{1}{n}y, i+1\right) \mu(dy)$$

converges to

$$\int e^{-\langle \alpha, y \rangle + H(\alpha)} V_F^n\left(x + \frac{1}{n}y, i+1\right) \mu(dy),$$

and thus the infimum may be restricted to \mathbb{Q}^d . ■

Proof of Lemma 2.2. Define

$$\underline{v} \doteq \sup_{\alpha \in \mathbb{R}^d} \inf_{\gamma \in \mathcal{C}} \left[\int f(y) \gamma(dy) + R(\gamma \|\mu) + \int \langle \alpha, y \rangle \gamma(dy) - H(\alpha) \right],$$

and similarly \bar{v} when the order of inf and sup are exchanged. The min/max inequality yields $\bar{v} \geq \underline{v}$. It suffices to show the reverse inequality.

For an arbitrary constant $M < \infty$, we have

$$\begin{aligned} \underline{v} &\geq \sup_{\|\alpha\| \leq M} \inf_{\gamma \in \mathcal{C}} \left[\int f(y) \gamma(dy) + R(\gamma \|\mu) + \int \langle \alpha, y \rangle \gamma(dy) - H(\alpha) \right] \\ &= \inf_{\gamma \in \mathcal{C}} \sup_{\|\alpha\| \leq M} \left[\int f(y) \gamma(dy) + R(\gamma \|\mu) + \int \langle \alpha, y \rangle \gamma(dy) - H(\alpha) \right] \\ &\doteq \underline{v}^M. \end{aligned}$$

The interchange of the infimum and supremum is valid thanks to the min/max theorem [38], the convexity of \mathcal{C} and $\{\alpha : \|\alpha\| \leq M\}$, and the compactness of $\{\alpha : \|\alpha\| \leq M\}$.

We will use the definition

$$L_M(\beta) \doteq \sup_{\|\alpha\| \leq M} [\langle \alpha, \beta \rangle - H(\alpha)].$$

Observe that $L_M(f y \mu(dy)) = 0$, since $0 \leq L_M \leq L$ and $L(f y \mu(dy)) = 0$. We also define the level set

$$\mathcal{D} \doteq \left\{ \gamma : \gamma \text{ is a probability measure on } \mathbb{R}^d, \text{ and } R(\gamma \|\mu) \leq 2\|f\|_\infty \right\} \subseteq \mathcal{C},$$

which is independent of M . We can write

$$\underline{v}^M = \inf_{\gamma \in \mathcal{C}} \left[\int f(y) \gamma(dy) + R(\gamma \|\mu) + L_M \left(\int y \gamma(dy) \right) \right].$$

It follows that $\underline{v}^M \leq \|f\|_\infty$ by taking $\gamma = \mu$. Since L_M is non-negative and f is bounded,

$$\inf_{\gamma \in \mathcal{C} \setminus \mathcal{D}} \left[\int f(y) \gamma(dy) + R(\gamma \|\mu) + L_M \left(\int y \gamma(dy) \right) \right] \geq -\|f\|_\infty + 2\|f\|_\infty \geq \underline{v}^M,$$

which implies that the infimum over \mathcal{C} is the same as the infimum over \mathcal{D} . Thus

$$\underline{v}^M = \inf_{\gamma \in \mathcal{D}} \left[\int f(y) \gamma(dy) + R(\gamma \|\mu) + L_M \left(\int y \gamma(dy) \right) \right].$$

We will further argue that there exists a $\gamma_M^* \in \mathcal{D}$ that achieves the infimum. To this end, we will associate the space of probability measures with the τ -topology, which is the smallest topology under which the mapping

$$\gamma \mapsto \int h(y) \gamma(dy)$$

is continuous for every bounded and measurable function h . The level set \mathcal{D} is not only compact in the weak topology, it is also compact under the τ -topology [13, Section 9.3]. Suppose now $\{\gamma_M^m : m \geq 1\}$ is a minimizing sequence. The compactness of \mathcal{D} implies the existence of a subsequence, still denoted by $\{\gamma_M^m\}$, such that

$$\gamma_M^m \rightarrow \gamma_M^*, \quad \text{for some } \gamma_M^* \in \mathcal{D}.$$

However, since the mapping

$$\gamma \mapsto R(\gamma \|\mu)$$

is lower semicontinuous (since the τ -topology is finer than the weak-convergence topology), we have

$$R(\gamma_M^* \|\mu) \leq \liminf_m R(\gamma_M^m \|\mu).$$

Furthermore, we have

$$\int f(y) \gamma_M^m(dy) \rightarrow \int f(y) \gamma_M^*(dy),$$

from the definition of the τ -topology and the boundedness of f , while

$$L_M \left(\int y \gamma_M^m(dy) \right) \rightarrow L_M \left(\int y \gamma_M^*(dy) \right),$$

thanks to the uniform integrability of $\{\gamma_M^m\}$ [13, Proposition 5.3.2] and the continuity of L_M . It follows readily that γ_M^* is a minimizer.

Again, thanks to the compactness of \mathcal{D} , there exists a subsequence of $\{\gamma_M^*\}$, still denoted by $\{\gamma_M^*\}$, such that $\gamma_M^* \rightarrow \gamma^* \in \mathcal{D}$ under the τ -topology, for some $\gamma^* \in \mathcal{D}$. Fix an arbitrary positive constant ε . Thanks to the lower-semicontinuity of $R(\cdot \|\mu)$, the definition of the τ -topology, and the boundedness of f , there exists $M_1 > 0$ such that for all $M \geq M_1$,

$$R(\gamma_M^* \|\mu) - R(\gamma^* \|\mu) \geq -\varepsilon,$$

$$\int f(y) \gamma_M^*(dy) - \int f(y) \gamma^*(dy) \geq -\varepsilon.$$

Furthermore, by the definition of L , there exists an $\alpha^* \in \mathbb{R}^d$ such that

$$L \left(\int y \gamma^*(dy) \right) \leq \left\langle \alpha^*, \int y \gamma^*(dy) \right\rangle - H(\alpha^*) + \varepsilon.$$

Since $\int y \gamma_M^*(dy) \rightarrow \int y \gamma^*(dy)$ thanks to the uniform integrability of $\{\gamma_M^*\}$, there also exists M_2 such that

$$\left\langle \alpha^*, \int y \gamma_M^*(dy) \right\rangle \geq \left\langle \alpha^*, \int y \gamma^*(dy) \right\rangle - \varepsilon,$$

for all $M \geq M_2$. Using the definition of L_M as a Legendre transform, for $M \geq \max\{\|\alpha^*\|, M_1, M_2\}$ we have

$$\begin{aligned} \underline{v} &\geq \underline{v}^M \\ &= \int f(y) \gamma_M^*(dy) + R(\gamma_M^* \|\mu) + L_M \left(\int y \gamma_M^*(dy) \right) \end{aligned}$$

$$\begin{aligned}
&\geq \int f(y)\gamma^*(dy) + R(\gamma^*\|\mu) + L_M \left(\int y\gamma_M^*(dy) \right) - 2\varepsilon \\
&\geq \int f(y)\gamma^*(dy) + R(\gamma^*\|\mu) + \left\langle \alpha^*, \int y\gamma_M^*(dy) \right\rangle - H(\alpha^*) - 2\varepsilon \\
&\geq \int f(y)\gamma^*(dy) + R(\gamma^*\|\mu) + \left\langle \alpha^*, \int y\gamma^*(dy) \right\rangle - H(\alpha^*) - 3\varepsilon \\
&\geq \int f(y)\gamma^*(dy) + R(\gamma^*\|\mu) + L \left(\int y\gamma^*(dy) \right) - 4\varepsilon \\
&\geq \inf_{\gamma \in \mathcal{C}} \left[\int f(y)\gamma(dy) + R(\gamma\|\mu) + L \left(\int y\gamma(dy) \right) \right] - 4\varepsilon.
\end{aligned}$$

Since ε is arbitrary, we obtain

$$\begin{aligned}
\underline{v} &\geq \inf_{\gamma \in \mathcal{C}} \left[\int f(y)\gamma(dy) + R(\gamma\|\mu) + L \left(\int y\gamma(dy) \right) \right] \\
&= \inf_{\gamma \in \mathcal{C}} \sup_{\alpha \in \mathbb{R}^d} \left[\int f(y)\gamma(dy) + R(\gamma\|\mu) + \int \langle \alpha, y \rangle \gamma(dy) - H(\alpha) \right] \\
&= \bar{v}.
\end{aligned}$$

This completes the proof. ■

References

- [1] S. Asmussen. Conjugate processes and the simulation of ruin probability. *Stochastic Process. Appl.*, 20:213–229, 1985.
- [2] S. Asmussen. Risk theory in a Markovian environment. *Scand. Actuarial J.*, pages 69–100, 1989.
- [3] S. Asmussen, R. Rubinstein, and C.L. Wang. Regenerative rare event simulation via likelihood ratios. *J. Appl. Probab.*, 31:797–815, 1994.
- [4] M. Bardi and I. Capuzzo-Dolcetta. *Optimal Control and Viscosity Solutions of Hamilton-Jacobi-Bellman Equations*. Birkhäuser, 1997.
- [5] D. Bertsekas and Shreve S. *Stochastic Optimal Control: The Discrete Time Case*. Academic Press, San Diego, California, 1978.
- [6] L. Breiman. *Probability Theory*. Addison-Wesley, Reading, Mass., 1968.
- [7] J. A. Bucklew. *Large Deviations Techniques in Decision, Simulation and Estimation*. Wiley, New York, 1990.

- [8] J.A. Bucklew. The blind simulation problem and regenerative processes. *IEEE Trans. Inform. Theory*, 44:2877–2891, 1998.
- [9] J.A. Bucklew, A. Ney, and J.S. Sadowsky. Monte Carlo simulation and large deviations theory for uniformly recurrent Markov chains. *J. Appl. Ann. Probab.*, 27:44–59, 1990.
- [10] J. Chen, D. Lu, J.S. Sadowsky, and K. Yao. On importance sampling in digital communications. Part I: Fundamentals. *IEEE J. Selected Areas in. Comm.*, 11:289–307, 1993.
- [11] J. F. Collamore. Importance sampling techniques for the multidimensional ruin problem for general Markov additive sequences of random vectors. *Ann. Appl. Prob.*, 12:382–421, 2002.
- [12] M. Cottrel, J.C. Fort, and G. Malgouyres. Large deviations and rare events in the study of stochastic algorithms. *IEEE Trans Automat. Control*, AC-28:907–920, 1983.
- [13] P. Dupuis and R. S. Ellis. *A Weak Convergence Approach to the Theory of Large Deviations*. John Wiley & Sons, New York, 1997.
- [14] P. Dupuis and H. J. Kushner. Stochastic systems with small noise, analysis and simulation; a phase locked loop example. *SIAM J. Appl. Math.*, 47:643–661, 1987.
- [15] P. Dupuis and H. J. Kushner. Minimizing exit probabilities; a large deviations approach. *SIAM J. on Control and Optimization*, 27:432–445, 1989.
- [16] P. Dupuis and A. Szpiro. Convergence of the optimal feedback policies in a numerical method for a class of deterministic optimal control problems. *SIAM J. on Control and Opt.*, 40:393–420, 2001.
- [17] R. J. Elliott and N. J. Kalton. *The Existence of Value in Differential Games*, volume 126 of *Memoirs of the Amer. Math. Society*. AMS, 1972.
- [18] J. Feng and T. Kurtz. Large deviation for stochastic processes. Technical report, The University of Massachusetts at Amherst, 2001.
- [19] W. H. Fleming. Stochastic control for small noise intensities. *SIAM J. on Control and Opt.*, 9:473–517, 1971.

- [20] W. H. Fleming. Exit probabilities and optimal stochastic control. *Applied Math. and Optimization*, 4:329–346, 1978.
- [21] W. H. Fleming and H. M. Soner. Asymptotic expansions for Markov processes with Levy generators. *Appl. Math. Optimization*, 19:203–223, 1989.
- [22] W. H. Fleming and P. E. Souganidis. Asymptotic series and the method of vanishing viscosity. *Indiana Univ. Math. J.*, 35:425–447, 1987.
- [23] P. Glasserman and S. Kou. Analysis of an importance sampling estimator for tandem queues. *ACM Trans. Modeling Comp. Simulation*, 4:22–42, 1995.
- [24] P. Glasserman and Y. Wang. Counter examples in importance sampling for large deviations probabilities. *Ann. Appl. Prob.*, 7:731–746, 1997.
- [25] P.W. Glynn. Large deviations for the infinite server queue in heavy traffic. *IMA Vol. Math. Appl.*, 71:387–394, 1995.
- [26] P. Heidelberger. Fast simulation of rare events in queueing and reliability models. *ACM Trans. Modeling Comp. Simulation*, 4:43–85, 1995.
- [27] H. J. Kushner and P. Dupuis. *Numerical Methods for Stochastic Control Problems in Continuous Time*. Springer-Verlag, New York, 2001. Revised Second Edition.
- [28] T. Lehtonen and H. Nyhrinen. Simulating level crossing probabilities by importance sampling. *Adv. Appl. Probab.*, 24:858–874, 1992a.
- [29] T. Lehtonen and H. Nyhrinen. On asymptotically efficient simulating of ruin probabilities in a Markovian environment. *Scand. Actuarial. J.*, 1:60–75, 1992b.
- [30] R.T. Rockafellar. *Convex Analysis*. Princeton University Press, Princeton, 1970.
- [31] J.S. Sadowsky. Large deviations and efficient simulation of excessive backlogs in a $GI/G/m$ queue. *IEEE. Trans. Automat. Control*, 36:1383–1394, 1991.
- [32] J.S. Sadowsky. On the optimality and stability of exponential twisting in Monte Carlo estimation. *IEEE. Trans. Inform. Theory*, 39:119–128, 1993.

- [33] J.S. Sadowsky. On Monte Carlo estimation of large deviations probabilities. *Ann. Appl. Prob.*, 6:399–422, 1996.
- [34] J.S. Sadowsky and J.A. Bucklew. On large deviations theory and asymptotically efficient Monte Carlo estimation. *IEEE Trans. Inform. Theory.*, 36:579–588, 1990.
- [35] P. Shahsbuddin. Importance sampling for the simulation of highly reliable Markovian systems. *Management Sciences*, 40:333–352, 1994.
- [36] S. J. Sheu. Stochastic control and exit probabilities of jump processes. *SIAM J. Control Optimization*, 23:306–328, 1985.
- [37] D. Siegmund. Importance sampling in the Monte Carlo study of sequential tests. *Ann. Statist.*, 4:673–684, 1976.
- [38] M. Sion. On general minimax theorems. *Pacific Journal of Mathematics*, 8:171–176, 1958.
- [39] S.R.S. Varadhan. *Large Deviations and Applications*. CBMS-NSF Regional Conference Series in Mathematics. SIAM, Philadelphia, 1984.