

Multilevel Monte Carlo Estimation of the Expected Value of Sample Information*

Tomohiko Hironaka[†], Michael B. Giles[‡], Takashi Goda[†], and Howard Thom[§]

Abstract. We study Monte Carlo estimation of the expected value of sample information (EVSI), which measures the expected benefit of gaining additional information for decision making under uncertainty. EVSI is defined as a nested expectation in which an outer expectation is taken with respect to one random variable Y and an inner conditional expectation with respect to the other random variable θ . Although the nested (Markov chain) Monte Carlo estimator has been often used in this context, a root-mean-square accuracy of ε is achieved notoriously at a cost of $O(\varepsilon^{-2-1/\alpha})$, where α denotes the order of convergence of the bias and is typically between $1/2$ and 1 . In this article we propose a novel efficient Monte Carlo estimator of EVSI by applying a multilevel Monte Carlo (MLMC) method. Instead of fixing the number of inner samples for θ as done in the nested Monte Carlo estimator, we consider a geometric progression on the number of inner samples, which yields a hierarchy of estimators on the inner conditional expectation with increasing approximation levels. Based on an elementary telescoping sum, our MLMC estimator is given by a sum of the Monte Carlo estimates of the differences between successive approximation levels on the inner conditional expectation. We show, under a set of assumptions on decision and information models, that successive approximation levels are tightly coupled, which directly proves that our MLMC estimator improves the necessary computational cost to optimal $O(\varepsilon^{-2})$. Numerical experiments confirm the considerable computational savings as compared to the nested Monte Carlo estimator.

Key words. expected value of sample information, multilevel Monte Carlo, nested expectations, decision-making under uncertainty

AMS subject classifications. 65C05, 62P10, 90B50

DOI. 10.1137/19M1284981

1. Introduction. Motivated by applications to medical decision making under uncertainty [27], we study Monte Carlo estimation of the expected value of sample information (EVSI). Let θ be a vector of random variables representing the uncertainty in the effectiveness of different medical treatments. Let D be a finite set of possible medical treatments, and for each treatment $d \in D$, f_d denotes a function of θ representing some measure of the patient outcome with “the larger the better,” where quality-adjusted life years (QALY) is typically employed in the context of medical decision making [1, 2, 25, 14]. Without any knowledge about θ , the best treatment is the one which maximizes the expectation of f_d , giving the

*Received by the editors September 3, 2019; accepted for publication (in revised form) July 15, 2020; published electronically September 30, 2020.

<https://doi.org/10.1137/19M1284981>

Funding: The work of the fourth author was supported by Medical Research Council grant MR/S036709/1.

[†]School of Engineering, University of Tokyo, Tokyo, Japan (hironaka-tomohiko@g.ecc.u-tokyo.ac.jp, goda@frcer.t.u-tokyo.ac.jp).

[‡]Mathematical Institute, University of Oxford, Oxford, UK (mike.giles@maths.ox.ac.uk).

[§]Bristol Medical School, University of Bristol, Bristol, UK (howard.thom@bristol.ac.uk).

average outcome

$$(1.1) \quad \max_{d \in D} \mathbb{E}_\theta [f_d(\theta)],$$

where $\mathbb{E}_\theta[\cdot]$ denotes the expectation taken with respect to the prior probability density function of θ . On the other hand, if perfect information on θ is available, the best treatment after knowing the value of θ is simply the one which maximizes $f_d(\theta)$, so that, on average, the outcome will be

$$\mathbb{E}_\theta \left[\max_{d \in D} f_d(\theta) \right].$$

The difference between these two values is called the expected value of perfect information (EVPI):

$$\text{EVPI} := \mathbb{E}_\theta \left[\max_{d \in D} f_d(\theta) \right] - \max_{d \in D} \mathbb{E}_\theta [f_d(\theta)].$$

However, it will be rare that we have access to perfect information on θ . In practice, what we obtain, for instance, through carrying out some new medical research is either *partial perfect information* or *sample information* on θ .

Partial perfect information is nothing but perfect information on only a subset of random variables θ_1 for a partition $\theta = (\theta_1, \theta_2)$, where θ_1 and θ_2 are assumed independent. After knowing the value of θ_1 , the best treatment is the one which maximizes the partial expectation $\mathbb{E}_{\theta_2} [f_d(\theta_1, \theta_2)]$. Therefore, the average outcome with partial perfect information on θ_1 will be

$$\mathbb{E}_{\theta_1} \left[\max_{d \in D} \mathbb{E}_{\theta_2} [f_d(\theta_1, \theta_2)] \right],$$

and the increment from (1.1) is called the expected value of partial perfect information (EVPPI):

$$\text{EVPPI} := \mathbb{E}_{\theta_1} \left[\max_{d \in D} \mathbb{E}_{\theta_2} [f_d(\theta_1, \theta_2)] \right] - \max_{d \in D} \mathbb{E}_\theta [f_d(\theta)].$$

Sample information on θ , which is of our interest in this article, is a single realization drawn from some probability distribution. To be more precise, we consider that information Y is stochastically generated according to the forward information model:

$$(1.2) \quad Y = h(\theta) + \epsilon,$$

where h is a known deterministic function of θ possibly with multiple outputs and ϵ is a zero-mean random variable with density ρ . Note that h and ϵ are called the observation operator and the observation noise, respectively [26, section 2]. It is widely known that Bayes' theorem provides an update of the probability density of θ after observing Y :

$$(1.3) \quad \pi^Y(\theta) = \frac{\rho(Y | \theta) \pi_0(\theta)}{\mathbb{E}_\theta[\rho(Y | \theta)]},$$

where $\pi_0(\theta)$ denotes the prior probability density of θ , and $\rho(Y | \theta)$ denotes the conditional probability density of Y given θ . Here $\rho(Y | \theta)$ is also called the likelihood of the information Y , and it follows from the model (1.2) that $\rho(Y | \theta) := \rho(Y - h(\theta))$.

Now, if such sample information Y is available, by choosing the best treatment which maximizes the conditional expectation $\mathbb{E}_{\theta|Y}[f_d(\theta)]$ depending on Y , where $\mathbb{E}_{\theta|Y}[\cdot]$ denotes the expectation taken with respect to the conditional probability density $\pi^Y(\theta)$, the overall average outcome becomes

$$\mathbb{E}_Y \left[\max_{d \in D} \mathbb{E}_{\theta|Y} [f_d(\theta)] \right].$$

Then EVSI represents the expected benefit of gaining the information Y and is defined by the difference

$$\text{EVSI} := \mathbb{E}_Y \left[\max_{d \in D} \mathbb{E}_{\theta|Y} [f_d(\theta)] \right] - \max_{d \in D} \mathbb{E}_{\theta} [f_d(\theta)].$$

In this article we are concerned with Monte Carlo estimation of EVSI. Given that EVPI can be estimated with root-mean-square accuracy ε by using $N = O(\varepsilon^{-2})$ independent and identically distributed (i.i.d.) samples of θ , denoted by $\theta^{(1)}, \dots, \theta^{(N)}$, as

$$\frac{1}{N} \sum_{n=1}^N \max_{d \in D} f_d(\theta^{(n)}) - \max_{d \in D} \frac{1}{N} \sum_{n=1}^N f_d(\theta^{(n)}),$$

it suffices to efficiently estimate the difference between EVPI and EVSI:

$$(1.4) \quad \text{EVPI} - \text{EVSI} = \mathbb{E}_{\theta} \left[\max_{d \in D} f_d(\theta) \right] - \mathbb{E}_Y \left[\max_{d \in D} \mathbb{E}_{\theta|Y} [f_d(\theta)] \right].$$

Because of the noncommutativity between the operators \mathbb{E} and $\max_{d \in D}$, this estimation is inherently a nested expectation problem, and it is far from trivial whether we can construct a good Monte Carlo estimator which achieves a root-mean-square accuracy ε at a cost of $O(\varepsilon^{-2})$.

Classically the most standard approach is to apply nested (Markov chain) Monte Carlo methods. For $M, N \in \mathbb{Z}_{>0}$, let $Y^{(1)}, \dots, Y^{(N)}$ be N outer i.i.d. samples of Y , and for each $1 \leq n \leq N$, let $\theta^{(n,1)}, \dots, \theta^{(n,M)}$ be M inner i.i.d. samples of θ conditional on $Y^{(n)}$. Then the nested Monte Carlo estimator of $\text{EVPI} - \text{EVSI}$ is given by

$$(1.5) \quad \frac{1}{N} \sum_{n=1}^N \left[\frac{1}{M} \sum_{m=1}^M \max_{d \in D} f_d(\theta^{(n,m)}) - \max_{d \in D} \frac{1}{M} \sum_{m=1}^M f_d(\theta^{(n,m)}) \right].$$

Here it is often hard to generate inner i.i.d. samples of θ conditional on some value of Y directly (although, conversely, it is quite easy to generate i.i.d. samples of Y conditional on some value of θ according to (1.2)). This is a major difference from estimating EVPPI.

To work around this difficulty, although the resulting samples are no longer i.i.d., one relies on Markov chain Monte Carlo (MCMC) sampling techniques such as Metropolis–Hastings sampling and Gibbs sampling; see [18, 20]. Under certain conditions, it follows from [16, 17] that one can establish a nonasymptotic error bound of $O(M^{-1/2})$ for MCMC estimation of the inner conditional expectation. Still, as inferred from a recent work of Giles and Goda [11] on EVPPI estimation, we need $N = O(\varepsilon^{-2})$ and $M = O(\varepsilon^{-1/\alpha})$ samples for outer and inner expectations, respectively, to estimate $\text{EVPI} - \text{EVSI}$ with root-mean-square accuracy ε . Here

α denotes the order of convergence of the bias and is typically between $1/2$ and 1 . This way the necessary total computational cost is of $O(\varepsilon^{-2-1/\alpha})$.

In this article, building upon the earlier work by Giles and Goda [11], we develop a novel efficient Monte Carlo estimator of EVPI–EVSI by using a multilevel Monte Carlo (MLMC) method [8, 9]. Although there has been extensive recent research on efficient approximations of EVSI in the medical decision making context [25, 14, 19, 15], our proposal avoids function approximations on the inner conditional expectation and any reliance on assumptions of multilinearity of f_d or weak correlation between random variables in θ . Recently MLMC estimators have been studied intensively for nested expectations of different forms, for instance, by [4, 12, 13]. We also refer the reader to [10] for a review of recent developments of MLMC applied to nested expectation problems. Importantly, our approach developed in this article does not require MCMC sampling for generating inner conditional samples of θ and can achieve a root-mean-square accuracy ε at a cost of optimal $O(\varepsilon^{-2})$. Moreover, it is straightforward to incorporate importance sampling techniques within our estimator, which may sometimes reduce the variance of the estimator significantly.

2. Multilevel Monte Carlo.

2.1. Basic theory. Before introducing our estimator of EVPI–EVSI, we give an overview of the MLMC method. Let P be a real-valued random variable which cannot be sampled exactly, and let P_0, P_1, \dots be a sequence of real-valued random variables which approximate P with increasing accuracy but also with increasing cost. In order to estimate $\mathbb{E}[P]$, we first approximate $\mathbb{E}[P]$ by $\mathbb{E}[P_L]$ for some $L \in \mathbb{Z}_{\geq 0}$ and then the standard Monte Carlo method estimates $\mathbb{E}[P_L]$ by using i.i.d. samples $P_L^{(1)}, P_L^{(2)}, \dots$ of P_L as

$$\mathbb{E}[P] \approx \mathbb{E}[P_L] \approx \overline{P}_L^N := \frac{1}{N} \sum_{n=1}^N P_L^{(n)}.$$

On the other hand, the MLMC method exploits the following telescoping sum representation:

$$\mathbb{E}[P_L] = \mathbb{E}[P_0] + \sum_{\ell=1}^L \mathbb{E}[P_\ell - P_{\ell-1}].$$

More generally, given a sequence of random variables $\Delta P_0, \Delta P_1, \dots$ which satisfy

$$\mathbb{E}[\Delta P_0] = \mathbb{E}[P_0] \quad \text{and} \quad \mathbb{E}[\Delta P_\ell] = \mathbb{E}[P_\ell - P_{\ell-1}] \quad \text{for } \ell \geq 1,$$

we have

$$\mathbb{E}[P_L] = \sum_{\ell=0}^L \mathbb{E}[\Delta P_\ell].$$

Then the MLMC estimator is given by a sum of independent Monte Carlo estimates of $\mathbb{E}[\Delta P_0], \mathbb{E}[\Delta P_1], \dots$, i.e.,

$$(2.1) \quad Z_{\text{MLMC}} = \sum_{\ell=0}^L \overline{\Delta P}_\ell^{N_\ell} = \sum_{\ell=0}^L \frac{1}{N_\ell} \sum_{n=1}^{N_\ell} \Delta P_\ell^{(n)}.$$

Since P_0, P_1, \dots approximate P with increasing accuracy, through a tight coupling of $P_{\ell-1}$ and P_ℓ , the variance of the correction variable ΔP_ℓ is expected to get smaller as the level ℓ increases. This implies that the numbers of samples N_0, N_1, \dots can also get smaller as the level ℓ increases so as to estimate each quantity $\mathbb{E}[\Delta P_\ell]$ accurately. If this is the case, the total computational cost can be reduced significantly as compared to the standard Monte Carlo method.

The following basic theorem from [8, 5, 9] makes the above observation explicit.

Theorem 2.1. *Let P be a random variable, and for $\ell \in \mathbb{Z}_{\geq 0}$, let P_ℓ be the level ℓ approximation of P . Assume that there exist independent correction random variables ΔP_ℓ with expected cost C_ℓ and variance V_ℓ , and positive constants $\alpha, \beta, \gamma, c_1, c_2, c_3$ such that $\alpha \geq \min(\beta, \gamma)/2$ and*

1. $\mathbb{E}[\Delta P_\ell] = \begin{cases} \mathbb{E}[P_0], & \ell = 0, \\ \mathbb{E}[P_\ell - P_{\ell-1}], & \ell \geq 1, \end{cases}$
2. $|\mathbb{E}[P_\ell - P]| \leq c_1 2^{-\alpha \ell},$
3. $V_\ell \leq c_2 2^{-\beta \ell},$
4. $C_\ell \leq c_3 2^{\gamma \ell}.$

Then there exists a positive constant c_4 such that, for any root-mean-square accuracy $\varepsilon < e^{-1}$, there are L and N_0, \dots, N_L for which the MLMC estimator (2.1) achieves a mean-square error less than ε^2 , i.e.,

$$\mathbb{E}[(Z_{\text{MLMC}} - \mathbb{E}[P])^2] \leq \varepsilon^2$$

with a computational cost C bounded above by

$$\mathbb{E}[C] \leq \begin{cases} c_4 \varepsilon^{-2}, & \beta > \gamma, \\ c_4 \varepsilon^{-2} (\log \varepsilon^{-1})^2, & \beta = \gamma, \\ c_4 \varepsilon^{-2 - (\gamma - \beta)/\alpha}, & \beta < \gamma. \end{cases}$$

Remark 2.2. As discussed in [5, section 2.1] and [11, section 2.1], under assumptions similar to those in Theorem 2.1, the standard Monte Carlo estimator achieves a root-mean-square accuracy at a cost of $O(\varepsilon^{-2 - \gamma/\alpha})$. Therefore, regardless of the values of $\beta > 0$ and γ , the MLMC estimator always has an asymptotically lower complexity bound than the standard Monte Carlo estimator.

2.2. MLMC estimator. Here we construct an MLMC estimator of the difference EVPI – EVSI. Our starting point is to insert (1.3) into (1.4), which results in

$$\begin{aligned} \text{EVPI} - \text{EVSI} &= \mathbb{E}_Y \mathbb{E}_{\theta|Y} \left[\max_{d \in D} f_d(\theta) \right] - \mathbb{E}_Y \left[\max_{d \in D} \mathbb{E}_{\theta|Y} [f_d(\theta)] \right] \\ &= \mathbb{E}_Y \left[\frac{\mathbb{E}_\theta [\max_{d \in D} f_d(\theta) \rho(Y|\theta)]}{\mathbb{E}_\theta [\rho(Y|\theta)]} - \max_{d \in D} \frac{\mathbb{E}_\theta [f_d(\theta) \rho(Y|\theta)]}{\mathbb{E}_\theta [\rho(Y|\theta)]} \right]. \end{aligned}$$

This has converted the posterior expectation with respect to θ given Y into the ratio of two prior expectations with respect to θ , which avoids the need for MCMC sampling. This idea has been used not only in the current context [19] but also in other areas related to Bayesian computations; see [22, 21, 6, 7], among many others. On a technical level, this gives a decisive difference from estimating EVPPI as considered in [11] for which we do not need such

treatment. As discussed later in subsection 3.2, efficient estimation of the nested expectation involving the ratio of two expectations is far from trivial. Although the algorithmic idea may perhaps seem a natural extension from [11], the numerical analysis of it is certainly not.

Based on the above expression, within the framework of the MLMC method, let us consider a real-valued random variable

$$P = \frac{\mathbb{E}_\theta [\max_{d \in D} f_d(\theta) \rho(Y | \theta)]}{\mathbb{E}_\theta [\rho(Y | \theta)]} - \max_{d \in D} \frac{\mathbb{E}_\theta [f_d(\theta) \rho(Y | \theta)]}{\mathbb{E}_\theta [\rho(Y | \theta)]},$$

with Y being the underlying random variable. We see that $\mathbb{E}_Y[P] = \text{EVPI} - \text{EVSI}$, but P cannot be sampled exactly because of the inner expectations.

It is important, however, that all these inner expectations are taken with respect to the prior probability density of θ , so that they can be approximated by the standard Monte Carlo method without requiring any MCMC sampling. This way, as a sequence of random variables P_0, P_1, \dots which approximate P with increasing accuracy, we consider the standard Monte Carlo estimation of P :

$$P_\ell := \frac{\overline{\max_{d \in D} f_d(\cdot) \rho(Y | \cdot)}^{M_\ell}}{\overline{\rho(Y | \cdot)}^{M_\ell}} - \max_{d \in D} \frac{\overline{f_d(\cdot) \rho(Y | \cdot)}^{M_\ell}}{\overline{\rho(Y | \cdot)}^{M_\ell}}$$

for an increasing sequence $M_0 < M_1 < \dots$ with $M_\ell \rightarrow \infty$ as $\ell \rightarrow \infty$. Here, in the definition of P_ℓ , we can use either the same M_ℓ samples of θ for all of the averages or independent M_ℓ samples of θ for each average. In this article we focus on the former approach and consider a geometric progression for M_0, M_1, \dots , i.e., let $M_\ell = M_0 2^\ell$ for some $M_0 \in \mathbb{Z}_{>0}$. In fact, the subsequent argument holds for a general integer base $b \geq 2$, but we restrict ourselves to the case $b = 2$ for simplicity of presentation.

Regarding a sequence of the correction variables $\Delta P_0, \Delta P_1, \dots$, following the ideas of [4, 11, 12], we consider an antithetic coupling of $P_{\ell-1}$ and P_ℓ . That is, the set of $M_0 2^\ell$ samples of θ used to compute P_ℓ is split into two disjoint sets of $M_0 2^{\ell-1}$ samples to compute two independent realizations of $P_{\ell-1}$, denoted by $P_{\ell-1}^{(a)}$ and $P_{\ell-1}^{(b)}$. Then ΔP_ℓ is defined by $\Delta P_0 = P_0$ and

$$\begin{aligned} \Delta P_\ell &:= P_\ell - \frac{P_{\ell-1}^{(a)} + P_{\ell-1}^{(b)}}{2} \\ &= \frac{\overline{\max_{d \in D} f_d(\cdot) \rho(Y | \cdot)}}{\overline{\rho(Y | \cdot)}} - \max_{d \in D} \frac{\overline{f_d(\cdot) \rho(Y | \cdot)}}{\overline{\rho(Y | \cdot)}} \\ &\quad - \frac{1}{2} \left[\frac{\overline{\max_{d \in D} f_d(\cdot) \rho(Y | \cdot)}^{(a)}}{\overline{\rho(Y | \cdot)}^{(a)}} + \frac{\overline{\max_{d \in D} f_d(\cdot) \rho(Y | \cdot)}^{(b)}}{\overline{\rho(Y | \cdot)}^{(b)}} \right] \\ &\quad + \frac{1}{2} \left[\max_{d \in D} \frac{\overline{f_d(\cdot) \rho(Y | \cdot)}^{(a)}}{\overline{\rho(Y | \cdot)}^{(a)}} + \max_{d \in D} \frac{\overline{f_d(\cdot) \rho(Y | \cdot)}^{(b)}}{\overline{\rho(Y | \cdot)}^{(b)}} \right] \end{aligned}$$

for $\ell \geq 1$, where we have omitted the superscripts M_ℓ from the first and second terms, and the averages with the superscripts (a) and (b) are taken by using the first and second

$M_0 2^{\ell-1}$ samples of θ used to compute the first two terms, respectively. Assuming that each computation of $f_d(\theta)$, $\rho(Y|\theta)$ and $h(\theta)$ can be performed with unit cost, it is clear that $\gamma = 1$ in Theorem 2.1 and $\mathbb{E}[\Delta P_\ell] = \mathbb{E}[P_\ell - P_{\ell-1}]$ for $\ell \geq 1$ because of the independence of the samples.

By the word ‘‘antithetic’’ we mean that the following properties hold:

$$(2.2) \quad \overline{\rho(Y|\cdot)} = \frac{\overline{\rho(Y|\cdot)^{(a)}} + \overline{\rho(Y|\cdot)^{(b)}}}{2},$$

$$(2.3) \quad \overline{f_d(\cdot)\rho(Y|\cdot)} = \frac{\overline{f_d(\cdot)\rho(Y|\cdot)^{(a)}} + \overline{f_d(\cdot)\rho(Y|\cdot)^{(b)}}}{2} \quad \text{for all } d \in D,$$

$$(2.4) \quad \overline{\max_{d \in D} f_d(\cdot)\rho(Y|\cdot)} = \frac{\overline{\max_{d \in D} f_d(\cdot)\rho(Y|\cdot)^{(a)}} + \overline{\max_{d \in D} f_d(\cdot)\rho(Y|\cdot)^{(b)}}}{2}.$$

This is the key advantage of the antithetic correction as compared to the standard correction $P_\ell - P_{\ell-1}^{(a)}$.

2.3. Combination with importance sampling. In practical applications, it is often the case where the likelihood $\rho(Y|\theta)$ (as a function of θ for a fixed Y) is highly concentrated around some values of θ . If one uses the i.i.d. samples of θ following from the prior density $\pi_0(\theta)$ for estimating $\mathbb{E}_\theta[\rho(Y|\theta)]$ and $\mathbb{E}_\theta[f_d(\theta)\rho(Y|\theta)]$, most of the samples can be distributed outside such concentrated regions. As a result, these quantities are estimated as almost zero, yielding a numerical instability of the estimates. Regarding theoretical analyses on such concentrated posterior measures on θ , we refer the reader to some recent papers [23, 24].

As discussed in [13, section 3.4] for MLMC applied to another nested expectation problem, one can combine importance sampling techniques with our MLMC estimator to address this issue. Let $q^Y(\theta)$ be an importance distribution of θ , conditional on Y , which needs to satisfy $q^Y(\theta) > 0$ for any θ with $\pi_0(\theta) > 0$. Since we have

$$\mathbb{E}_\theta[f_d(\theta)\rho(Y|\theta)] = \mathbb{E}_{\theta \sim q^Y} \left[f_d(\theta) \frac{\rho(Y|\theta)\pi_0(\theta)}{q^Y(\theta)} \right]$$

and

$$\mathbb{E}_\theta[\rho(Y|\theta)] = \mathbb{E}_{\theta \sim q^Y} \left[\frac{\rho(Y|\theta)\pi_0(\theta)}{q^Y(\theta)} \right],$$

the random variables P_ℓ and ΔP_ℓ can be replaced, respectively, by

$$\begin{aligned} P_\ell &= \frac{\overline{\max_{d \in D} f_d(\cdot)\rho(Y|\cdot)\pi_0(\cdot)/q^Y(\cdot)}^{M_\ell}}{\overline{\rho(Y|\cdot)\pi_0(\cdot)/q^Y(\cdot)}^{M_\ell}} - \max_{d \in D} \frac{\overline{f_d(\cdot)\rho(Y|\cdot)\pi_0(\cdot)/q^Y(\cdot)}^{M_\ell}}{\overline{\rho(Y|\cdot)\pi_0(\cdot)/q^Y(\cdot)}^{M_\ell}}, \\ \Delta P_\ell &= \frac{\overline{\max_{d \in D} f_d(\cdot)\rho(Y|\cdot)\pi_0(\cdot)/q^Y(\cdot)}}{\overline{\rho(Y|\cdot)\pi_0(\cdot)/q^Y(\cdot)}} - \max_{d \in D} \frac{\overline{f_d(\cdot)\rho(Y|\cdot)\pi_0(\cdot)/q^Y(\cdot)}}{\overline{\rho(Y|\cdot)\pi_0(\cdot)/q^Y(\cdot)}} \\ &\quad - \frac{1}{2} \left[\frac{\overline{\max_{d \in D} f_d(\cdot)\rho(Y|\cdot)\pi_0(\cdot)/q^Y(\cdot)^{(a)}}}{\overline{\rho(Y|\cdot)\pi_0(\cdot)/q^Y(\cdot)^{(a)}}} + \frac{\overline{\max_{d \in D} f_d(\cdot)\rho(Y|\cdot)\pi_0(\cdot)/q^Y(\cdot)^{(b)}}}{\overline{\rho(Y|\cdot)\pi_0(\cdot)/q^Y(\cdot)^{(b)}}} \right] \\ &\quad + \frac{1}{2} \left[\max_{d \in D} \frac{\overline{f_d(\cdot)\rho(Y|\cdot)\pi_0(\cdot)/q^Y(\cdot)^{(a)}}}{\overline{\rho(Y|\cdot)\pi_0(\cdot)/q^Y(\cdot)^{(a)}}} + \max_{d \in D} \frac{\overline{f_d(\cdot)\rho(Y|\cdot)\pi_0(\cdot)/q^Y(\cdot)^{(b)}}}{\overline{\rho(Y|\cdot)\pi_0(\cdot)/q^Y(\cdot)^{(b)}}} \right], \end{aligned}$$

where all of the averages are taken with respect to i.i.d. samples of $\theta \sim q^Y$ for a randomly chosen Y .

We suggest finding a good approximation of the posterior distribution $\pi^Y(\theta)$ for $q^Y(\theta)$. If one can do so, it follows from Bayes' theorem (1.3) that

$$\frac{\rho(Y|\theta)\pi_0(\theta)}{q^Y(\theta)} \approx \frac{\rho(Y|\theta)\pi_0(\theta)}{\pi^Y(\theta)} = \mathbb{E}_\theta[\rho(Y|\theta)].$$

Since the rightmost side does not depend on θ , the integrand $\rho(Y|\theta)\pi_0(\theta)/q^Y(\theta)$ appearing in the denominator of each term for P_ℓ and ΔP_ℓ becomes close to a constant function, so that its variance is extremely small. This way we can avoid the numerical instability of our original MLMC estimator.

3. Theoretical results. In this section, we prove $\beta > \gamma = 1$ and $\alpha \geq \beta/2$ under a set of assumptions on the decision and information models. This directly implies from Theorem 2.1 that our MLMC estimator of EVPI – EVSI achieves a root-mean-square accuracy ε at a cost of optimal $O(\varepsilon^{-2})$.

In what follows, for simplicity of notation, we write $\rho(Y) = \mathbb{E}_\theta[\rho(Y|\theta)]$ and $F_d(Y) = \mathbb{E}_\theta[f_d(\theta)\rho(Y|\theta)]$. Note that we have

$$\mathbb{E} \left[\frac{\overline{\rho(Y|\cdot)}}{\rho(Y)} \mid Y \right] = \mathbb{E}_\theta \left[\frac{\rho(Y|\theta)}{\rho(Y)} \mid Y \right] = 1.$$

Moreover, we write

$$\bar{g}_d^{(a)} = \frac{\overline{f_d(\cdot)\rho(Y|\cdot)^{(a)}}}{\overline{\rho(Y|\cdot)^{(a)}}}, \quad \bar{g}_d^{(b)} = \frac{\overline{f_d(\cdot)\rho(Y|\cdot)^{(b)}}}{\overline{\rho(Y|\cdot)^{(b)}}}, \quad \bar{g}_d = \frac{\overline{f_d(\cdot)\rho(Y|\cdot)}}{\overline{\rho(Y|\cdot)}}$$

for $d \in D$, and

$$\bar{g}_{\max}^{(a)} = \frac{\overline{\max_{d \in D} f_d(\cdot)\rho(Y|\cdot)^{(a)}}}{\overline{\rho(Y|\cdot)^{(a)}}}, \quad \bar{g}_{\max}^{(b)} = \frac{\overline{\max_{d \in D} f_d(\cdot)\rho(Y|\cdot)^{(b)}}}{\overline{\rho(Y|\cdot)^{(b)}}}, \quad \bar{g}_{\max} = \frac{\overline{\max_{d \in D} f_d(\cdot)\rho(Y|\cdot)}}{\overline{\rho(Y|\cdot)}}.$$

Then ΔP_ℓ is given by $\Delta P_\ell = \Delta P_{\ell,1} + \Delta P_{\ell,2}$ with

$$\Delta P_{\ell,1} = \bar{g}_{\max} - \frac{1}{2} \left(\bar{g}_{\max}^{(a)} + \bar{g}_{\max}^{(b)} \right) \quad \text{and} \quad \Delta P_{\ell,2} = \frac{1}{2} \left(\max_{d \in D} \bar{g}_d^{(a)} + \max_{d \in D} \bar{g}_d^{(b)} \right) - \max_{d \in D} \bar{g}_d.$$

For a given Y , we define

$$G_d(Y) := \mathbb{E}_{\theta|Y} [f_d(\theta)] = \frac{\mathbb{E}_\theta [f_d(\theta)\rho(Y|\theta)]}{\mathbb{E}_\theta [\rho(Y|\theta)]} = \frac{F_d(Y)}{\rho(Y)}$$

for each $d \in D$, and

$$d_{\text{opt}}(Y) := \arg \max_{d \in D} F_d(Y) = \arg \max_{d \in D} G_d(Y).$$

The second equality is trivial since the denominator of G_d does not affect the choice $\max_{d \in D}$. The domain for Y is divided into a number of subdomains in which the optimal decision d_{opt} is unique. We denote by K the dividing decision manifold on which d_{opt} is not uniquely defined, and we assume that K is a lower-dimensional subspace of the domain for Y .

Let us give some assumptions on the decision and information models.

Assumption 3.1. There exists a constant $0 < F_{\max} < \infty$ such that $|f_d(\theta)| \leq F_{\max}$ for any d and θ .

Assumption 3.2. There exists a constant $c_0 > 0$ such that for all $0 < \epsilon < 1$

$$\mathbb{P}_Y \left[\min_{y \in K} \|Y - y\| \leq \epsilon \right] \leq c_0 \epsilon.$$

Assumption 3.3. There exist constants $c_1, c_2 > 0$ such that if $Y \notin K$, the following holds:

$$\max_{d \in D} G_d(Y) - \max_{\substack{d \in D \\ d \neq d_{\text{opt}}(Y)}} G_d(Y) > \min \left(c_1, c_2 \min_{y \in K} \|Y - y\| \right).$$

Assumption 3.4. There exists a constant $p \geq 2$ such that

$$\mathbb{E}_Y \left[\mathbb{E}_\theta \left[\left(\frac{\rho(Y|\theta)}{\rho(Y)} \right)^p \right] \right] < \infty.$$

Remark 3.5. Assumptions 3.1–3.3 are similar to those considered in [11]. In particular, Assumption 3.2 is introduced to ensure a bound on the probability that Y is close to the decision manifold K , while Assumption 3.3 is to ensure a linear separation of different decisions as Y moves away from K . Assumption 3.1, which is stronger than [11, Assumption 1], together with Assumption 3.4 enables us to bound the difference between $\bar{g}_d^{(a)}$, $\bar{g}_d^{(b)}$, \bar{g}_d , and G_d .

Now we are ready to state our main result of this article.

Theorem 3.6. *If Assumptions 3.1–3.4 are satisfied, we have*

$$\mathbb{V}[\Delta P_\ell] = O(2^{-(3p/(2p+2))\ell}) \quad \text{and} \quad \mathbb{E}[|\Delta P_\ell|] = O(2^{-(p/(p+1))\ell}),$$

where $p \geq 2$ is as given in Assumption 3.4.

Remark 3.7. When an importance sampling is used within the MLMC estimator, the same orders of the variance and the mean of ΔP_ℓ can be shown by replacing Assumption 3.4 with the existence of a constant $p \geq 2$ such that

$$\mathbb{E}_Y \left[\mathbb{E}_\theta \left[\left(\frac{\rho(Y|\theta)\pi_0(\theta)}{\rho(Y)q^Y(\theta)} \right)^p \right] \right] < \infty.$$

Since the result can be proven in the same manner with the original MLMC estimator, we shall give a proof of Theorem 3.6 only for the original estimator.

This theorem implies that the parameters α and β in Theorem 2.1 are equal to $p/(p+1)$ and $3p/(2p+2)$, respectively. Since $\gamma = 1 < \beta$ if $p > 2$, our MLMC estimator of EVPI – EVSI is in the first regime. Therefore, the total computational complexity to achieve a root-mean-square accuracy ε is of order ε^{-2} . If $p = 2$, on the other hand, the equality $\beta = \gamma$ holds, which means that our MLMC estimator is in the second regime. In the next subsection, we give a proof of this theorem by using several lemmas which are shown later in subsection 3.2.

3.1. Proof of the main result. We follow an argument similar to that used in [11, Theorem 3] in conjunction with novel results shown later. Recalling that $\Delta P_\ell = \Delta P_{\ell,1} + \Delta P_{\ell,2}$, we have

$$\mathbb{V}[\Delta P_\ell] \leq \mathbb{E}[|\Delta P_\ell|^2] \leq 2\mathbb{E}[|\Delta P_{\ell,1}|^2] + 2\mathbb{E}[|\Delta P_{\ell,2}|^2].$$

We shall see later, in Remark 3.10, that the first term on the right-hand side is of $O(2^{-\min(p,4)\ell/2})$, which decays no slower than the desired order $2^{-(3p/(2p+2))\ell}$ for any $p \geq 2$. Thus it suffices to prove that the second term is of $O(2^{-(3p/(2p+2))\ell})$.

For p as given in Assumption 3.4, let us define $\epsilon = 2^{-(p/(2p+2))\ell}$ and consider the events

$$A \equiv \left\{ \min_{y \in K} \|Y - y\| \leq \epsilon \right\},$$

$$B \equiv \bigcup_{d \in D} \left\{ \max \left(|\bar{g}_d^{(a)} - G_d|, |\bar{g}_d^{(b)} - G_d|, |\bar{g}_d - G_d| \right) \geq \frac{1}{2} c_2 \epsilon \right\},$$

where c_2 is as defined in Assumption 3.3.

For an event E , let $\mathbf{1}_E$ denote the indicator function, which is 1 if $\omega \in E$ and zero otherwise, and let E^c denote the complement of E . By using Hölder's inequality, we have

$$\begin{aligned} \mathbb{E}[|\Delta P_{\ell,2}|^2] &= \mathbb{E}[|\Delta P_{\ell,2}|^2 \mathbf{1}_{A \cup B}] + \mathbb{E}[|\Delta P_{\ell,2}|^2 \mathbf{1}_{A^c \cap B^c}] \\ &\leq (\mathbb{E}[|\Delta P_{\ell,2}|^p])^{2/p} (\mathbb{E}[\mathbf{1}_{A \cup B}^{p/(p-2)}])^{(p-2)/p} + \mathbb{E}[|\Delta P_{\ell,2}|^2 \mathbf{1}_{A^c \cap B^c}] \\ &\leq (\mathbb{E}[|\Delta P_{\ell,2}|^p])^{2/p} (\mathbb{P}[A] + \mathbb{P}[B])^{(p-2)/p} + \mathbb{E}[|\Delta P_{\ell,2}|^2 \mathbf{1}_{A^c \cap B^c}]. \end{aligned}$$

In the following we show bounds on $\mathbb{E}[|\Delta P_{\ell,2}|^p]$, $\mathbb{P}[A] + \mathbb{P}[B]$, and $\mathbb{E}[|\Delta P_{\ell,2}|^2 \mathbf{1}_{A^c \cap B^c}]$, respectively.

Bound on $\mathbb{E}[|\Delta P_{\ell,2}|^p]$. Noting that the inequality

$$\left| \max_{d \in D} a_d - \max_{d \in D} b_d \right| \leq \max_{d \in D} |a_d - b_d| \leq \sum_{d \in D} |a_d - b_d|$$

holds for any two $|D|$ -dimensional vectors with component a_d, b_d and applying Jensen's inequality twice, we get

$$\begin{aligned} |\Delta P_{\ell,2}|^p &= \left| \frac{1}{2} \left(\max_{d \in D} \bar{g}_d^{(a)} + \max_{d \in D} \bar{g}_d^{(b)} \right) - \max_{d \in D} \bar{g}_d \right|^p \\ &= \left| \frac{1}{2} \left(\max_{d \in D} \bar{g}_d^{(a)} - \max_{d \in D} G_d \right) + \frac{1}{2} \left(\max_{d \in D} \bar{g}_d^{(b)} - \max_{d \in D} G_d \right) - \left(\max_{d \in D} \bar{g}_d - \max_{d \in D} G_d \right) \right|^p \\ &\leq \left| \sum_{d \in D} \left(\frac{1}{2} |\bar{g}_d^{(a)} - G_d| + \frac{1}{2} |\bar{g}_d^{(b)} - G_d| + |\bar{g}_d - G_d| \right) \right|^p \\ &\leq |D|^{p-1} \sum_{d \in D} \left(\frac{1}{2} |\bar{g}_d^{(a)} - G_d| + \frac{1}{2} |\bar{g}_d^{(b)} - G_d| + |\bar{g}_d - G_d| \right)^p \\ &\leq (2|D|)^{p-1} \sum_{d \in D} \left(\frac{1}{2} |\bar{g}_d^{(a)} - G_d|^p + \frac{1}{2} |\bar{g}_d^{(b)} - G_d|^p + |\bar{g}_d - G_d|^p \right). \end{aligned}$$

It follows from the first assertion of Lemma 3.9 that

$$\mathbb{E} \left[\left| \overline{g_d^{(a)}} - G_d \right|^p \right], \mathbb{E} \left[\left| \overline{g_d^{(b)}} - G_d \right|^p \right], \mathbb{E} \left[\left| \overline{g_d} - G_d \right|^p \right] = O(2^{-p\ell/2})$$

for any $d \in D$. This leads to a bound on $\mathbb{E}[|\Delta P_{\ell,2}|^p]$ of order $2^{-p\ell/2}$.

Bound on $\mathbb{P}[A] + \mathbb{P}[B]$. It is straightforward from Assumption 3.2 that $\mathbb{P}[A] \leq c_0\epsilon$. Regarding a bound on $\mathbb{P}[B]$, we have

$$\mathbb{P}[B] \leq \sum_{d \in D} \left(\mathbb{P} \left[\left| \overline{g_d^{(a)}} - G_d \right| \geq \frac{1}{2}c_2\epsilon \right] + \mathbb{P} \left[\left| \overline{g_d^{(b)}} - G_d \right| \geq \frac{1}{2}c_2\epsilon \right] + \mathbb{P} \left[\left| \overline{g_d} - G_d \right| \geq \frac{1}{2}c_2\epsilon \right] \right).$$

Using the Markov inequality, we obtain

$$\begin{aligned} \mathbb{P} \left[\left| \overline{g_d} - G_d \right| \geq \frac{1}{2}c_2\epsilon \mid Y \right] &= \mathbb{P} \left[\left| \overline{g_d} - G_d \right|^p \geq \frac{1}{2^p}c_2^p\epsilon^p \mid Y \right] \\ &\leq \frac{2^p}{c_2^p\epsilon^p} \mathbb{E} \left[\left| \overline{g_d} - G_d \right|^p \mid Y \right]. \end{aligned}$$

Taking an outer expectation with respect to Y , the tower property provides

$$\begin{aligned} \mathbb{P} \left[\left| \overline{g_d} - G_d \right| \geq \frac{1}{2}c_2\epsilon \right] &= \mathbb{E}_Y \left[\mathbb{P} \left[\left| \overline{g_d} - G_d \right| \geq \frac{1}{2}c_2\epsilon \mid Y \right] \right] \\ &\leq \frac{2^p}{c_2^p\epsilon^p} \mathbb{E} \left[\left| \overline{g_d} - G_d \right|^p \right]. \end{aligned}$$

The first assertion of Lemma 3.9 states that the expectation on the rightmost side above is of order $2^{-p\ell/2}$. Since similar bounds exist for $\mathbb{P} \left[\left| \overline{g_d^{(a)}} - G_d \right| \geq \frac{1}{2}c_2\epsilon \right]$ and $\mathbb{P} \left[\left| \overline{g_d^{(b)}} - G_d \right| \geq \frac{1}{2}c_2\epsilon \right]$, this proves a bound on $\mathbb{P}[A] + \mathbb{P}[B]$ of order

$$\max(\epsilon, \epsilon^{-p}2^{-p\ell/2}) = 2^{-(p/(2p+2))\ell}.$$

Bound on $\mathbb{E}[|\Delta P_{\ell,2}|^2 \mathbf{1}_{A^c \cap B^c}]$. Finally let us consider the case where the event $A^c \cap B^c$ happens. Namely we have $\min_{y \in K} \|Y - y\| > \epsilon$, and for any $d \in D$

$$\left| \overline{g_d^{(a)}} - G_d \right|, \left| \overline{g_d^{(b)}} - G_d \right|, \left| \overline{g_d} - G_d \right| < \frac{1}{2}c_2\epsilon.$$

For a particular value of Y , for any decision $d \neq d_{\text{opt}}(Y)$, Assumption 3.3 ensures

$$\begin{aligned} \overline{g_{d_{\text{opt}}}} - \overline{g_d} &= (G_{d_{\text{opt}}} - G_d) + (\overline{g_{d_{\text{opt}}}} - G_{d_{\text{opt}}}) - (\overline{g_d} - G_d) \\ &> \min(c_1, c_2\epsilon) - \frac{1}{2}c_2\epsilon - \frac{1}{2}c_2\epsilon = \min(c_1 - c_2\epsilon, 0). \end{aligned}$$

Therefore, for a sufficiently large ℓ , we have $c_1 - c_2\epsilon > 0$, and so $\overline{g_{d_{\text{opt}}}} - \overline{g_d} > 0$. This means that $\arg \max_{d \in D} \overline{g_d} = d_{\text{opt}}$, that is, we shall always choose an optimal decision, and the same for $\overline{g_d^{(a)}}$ and $\overline{g_d^{(b)}}$. It follows that

$$\Delta P_{\ell,2} = \frac{1}{2} \left(\max_{d \in D} \overline{g_d^{(a)}} + \max_{d \in D} \overline{g_d^{(b)}} \right) - \max_{d \in D} \overline{g_d} = \frac{1}{2} \left(\overline{g_{d_{\text{opt}}^{(a)}}} + \overline{g_{d_{\text{opt}}^{(b)}}} \right) - \overline{g_{d_{\text{opt}}}}.$$

We see from the second assertion of Lemma 3.9 with $q = 2$ that the expectation of the square of the rightmost side above is of order $2^{-\min(p,4)\ell/2}$, which proves a bound on $\mathbb{E}[|\Delta P_{\ell,2}|^2 \mathbf{1}_{A^c \cap B^c}]$.

Altogether, $\mathbb{V}[\Delta P_\ell]$ is upper bounded as

$$\begin{aligned} \mathbb{V}[\Delta P_\ell] &\leq 2\mathbb{E}[|\Delta P_{\ell,1}|^2] + 2\mathbb{E}[|\Delta P_{\ell,2}|^2] \\ &= O(2^{-\min(p,4)\ell/2}) + O(2^{-\ell}) \times O(2^{-(p-2)\ell/(2p+2)}) + O(2^{-\min(p,4)\ell/2}) \\ &= O(2^{-(3p/(2p+2))\ell}). \end{aligned}$$

Since the bound on $\mathbb{E}[|\Delta P_\ell|]$ can be proven in a similar way, we omit the proof.

3.2. Auxiliary results. Before stating necessary auxiliary results, let us first recall the following fact. The proof can be found in [11, Lemma 1].

Lemma 3.8. *Let X be a real-valued random variable with mean zero, and let \bar{X}_N be an average of N i.i.d. samples of X . If $\mathbb{E}[|X|^p] < \infty$ for $p \geq 2$, there exists a constant $C_p > 0$ depending only on p such that*

$$\mathbb{E}[|\bar{X}_N|^p] \leq C_p \frac{\mathbb{E}[|X|^p]}{N^{p/2}} \quad \text{and} \quad \mathbb{P}[|\bar{X}_N| > c] \leq C_p \frac{\mathbb{E}[|X|^p]}{c^p N^{p/2}}.$$

Lemma 3.9. *If Assumptions 3.1 and 3.4 are satisfied, then for any $d \in D$ we have*

$$\mathbb{E}[|\bar{g}_d - G_d|^q] = O(2^{-q\ell/2}) \quad \text{for any } 1 \leq q \leq p$$

and

$$\mathbb{E} \left[\left| \frac{1}{2} \left(\bar{g}_d^{(a)} + \bar{g}_d^{(b)} \right) - \bar{g}_d \right|^q \right] = O(2^{-\min(p,2q)\ell/2}) \quad \text{for any } 1 \leq q \leq p,$$

where p is as defined in Assumption 3.4.

Proof. Note that because of the fact that the same i.i.d. samples of θ are used both in the denominator and numerator of $\bar{g}_d^{(a)}$, $\bar{g}_d^{(b)}$, \bar{g}_d , respectively, it follows from Assumption 3.1 that $|\bar{g}_d^{(a)}|, |\bar{g}_d^{(b)}|, |\bar{g}_d| \leq F_{max}$. Also we have

$$|F_d(Y)| \leq F_{max} \mathbb{E}_\theta [\rho(Y | \theta)] = F_{max} \rho(Y).$$

For the first assertion, we give a proof only for $q = p$, since the result for $q < p$ follows immediately from Hölder's inequality.

We define an extreme event E as

$$E : \left| \frac{\rho(Y | \cdot)}{\rho(Y)} - 1 \right| > \frac{1}{2}$$

and then

$$\mathbb{E}[|\bar{g}_d - G_d|^p] = \mathbb{E}[|\bar{g}_d - G_d|^p \mathbf{1}_E] + \mathbb{E}[|\bar{g}_d - G_d|^p \mathbf{1}_{E^c}].$$

For the first term, due to Hölder's inequality, the tower property of expectation, and Lemma 3.8, we have

$$\begin{aligned} \mathbb{E}[|\bar{g}_d - G_d|^p \mathbf{1}_E] &\leq (2F_{max})^p \mathbb{P}[E] \\ (3.1) \quad &\leq 2^{2p-p\ell/2} M_0^{-p/2} C_p F_{max}^p \mathbb{E}_Y \left[\mathbb{E}_\theta \left[\left| \frac{\rho(Y | \theta)}{\rho(Y)} - 1 \right|^p \right] \right]. \end{aligned}$$

For the second term, when $|\overline{\rho(Y|\cdot)}/\rho(Y) - 1| \leq 1/2$ we have

$$\begin{aligned} |\overline{g_d} - G_d|^p &\leq 2^{p-1} \left| \frac{f_d(\cdot)\rho(Y|\cdot)}{\rho(Y|\cdot)} - \frac{F_d(Y)}{\rho(Y|\cdot)} \right|^p + 2^{p-1} \left| \frac{F_d(Y)}{\rho(Y|\cdot)} - \frac{F_d(Y)}{\rho(Y)} \right|^p \\ &\leq \frac{2^{2p-1}}{\rho(Y)^p} |f_d(\cdot)\rho(Y|\cdot) - F_d(Y)|^p + 2^{2p-1} F_{max}^p \left| \frac{\rho(Y|\cdot)}{\rho(Y)} - 1 \right|^p, \end{aligned}$$

and therefore, again by Lemma 3.8,

$$\begin{aligned} \mathbb{E}[|\overline{g_d} - G_d|^p \mathbf{1}_{E^c}] &\leq 2^{2p-1-p\ell/2} M_0^{-p/2} C_p \mathbb{E}_Y \left[\mathbb{E}_\theta \left[\left| \frac{f_d(\theta)\rho(Y|\theta) - F_d(Y)}{\rho(Y)} \right|^p \right] \right] \\ (3.2) \quad &+ 2^{2p-1-p\ell/2} M_0^{-p/2} C_p F_{max}^p \mathbb{E}_Y \left[\mathbb{E}_\theta \left[\left| \frac{\rho(Y|\theta)}{\rho(Y)} - 1 \right|^p \right] \right]. \end{aligned}$$

Note that

$$\left| \frac{\rho(Y|\theta)}{\rho(Y)} - 1 \right|^p \leq \max \left(\frac{\rho(Y|\theta)}{\rho(Y)}, 1 \right)^p \leq \left(\frac{\rho(Y|\theta)}{\rho(Y)} \right)^p + 1,$$

which gives

$$(3.3) \quad \mathbb{E}_Y \left[\mathbb{E}_\theta \left[\left| \frac{\rho(Y|\theta)}{\rho(Y)} - 1 \right|^p \right] \right] \leq \mathbb{E}_Y \left[\mathbb{E}_\theta \left[\left(\frac{\rho(Y|\theta)}{\rho(Y)} \right)^p \right] \right] + 1.$$

Similarly,

$$\begin{aligned} |f_d(\theta)\rho(Y|\theta) - F_d(Y)|^p &\leq 2^{p-1} (|f_d(\theta)\rho(Y|\theta)|^p + |F_d(Y)|^p) \\ &\leq 2^{p-1} F_{max}^p (\rho(Y|\theta)^p + \rho(Y)^p), \end{aligned}$$

which gives

$$(3.4) \quad \mathbb{E}_Y \left[\mathbb{E}_\theta \left[\left| \frac{f_d(\theta)\rho(Y|\theta) - F_d(Y)}{\rho(Y)} \right|^p \right] \right] \leq 2^{p-1} F_{max}^p \left(\mathbb{E}_Y \left[\mathbb{E}_\theta \left[\left(\frac{\rho(Y|\theta)}{\rho(Y)} \right)^p \right] \right] + 1 \right).$$

Combining the bounds in (3.1) and (3.2), and then using (3.3) and (3.4), we get

$$\mathbb{E}[|\overline{g_d} - G_d|^p] \leq (3 \cdot 2^{2p-1} + 2^{3p-2}) 2^{-p\ell/2} M_0^{-p/2} C_p F_{max}^p \left(\mathbb{E}_Y \left[\mathbb{E}_\theta \left[\left(\frac{\rho(Y|\theta)}{\rho(Y)} \right)^p \right] \right] + 1 \right),$$

which completes the proof of the first assertion.

For the second assertion we define the extreme event E as

$$E : \max \left(\left| \frac{\rho(Y|\cdot)^{(a)}}{\rho(Y)} - 1 \right|, \left| \frac{\rho(Y|\cdot)^{(b)}}{\rho(Y)} - 1 \right| \right) > \frac{1}{2}$$

and then

$$\begin{aligned} \mathbb{E} \left[\left| \frac{1}{2} (\overline{g_d}^{(a)} + \overline{g_d}^{(b)}) - \overline{g_d} \right|^q \right] &= \mathbb{E} \left[\left| \frac{1}{2} (\overline{g_d}^{(a)} + \overline{g_d}^{(b)}) - \overline{g_d} \right|^q \mathbf{1}_E \right] \\ &+ \mathbb{E} \left[\left| \frac{1}{2} (\overline{g_d}^{(a)} + \overline{g_d}^{(b)}) - \overline{g_d} \right|^q \mathbf{1}_{E^c} \right]. \end{aligned}$$

For the first term, following similar reasoning to before, we have

$$\begin{aligned}
 & \mathbb{E} \left[\left| \frac{1}{2} \left(\overline{g_d^{(a)}} + \overline{g_d^{(b)}} \right) - \overline{g_d} \right|^q \mathbf{1}_E \right] \\
 & \leq (2F_{max})^q \left(\mathbb{P} \left[\left| \frac{\overline{\rho(Y|\cdot)^{(a)}}}{\rho(Y)} - 1 \right| > \frac{1}{2} \right] + \mathbb{P} \left[\left| \frac{\overline{\rho(Y|\cdot)^{(b)}}}{\rho(Y)} - 1 \right| > \frac{1}{2} \right] \right) \\
 (3.5) \quad & \leq 2^{p+q+1} C_p 2^{-p(\ell-1)/2} M_0^{-p/2} F_{max}^q \left(\mathbb{E}_Y \left[\mathbb{E}_\theta \left[\left(\frac{\rho(Y|\theta)}{\rho(Y)} \right)^p \right] \right] + 1 \right).
 \end{aligned}$$

For the second term, we use the antithetic properties (2.2) and (2.3) to give

$$\begin{aligned}
 \frac{1}{2} \left(\overline{g_d^{(a)}} + \overline{g_d^{(b)}} \right) - \overline{g_d} &= \frac{1}{2} \left(\overline{f_d(\cdot)\rho(Y|\cdot)^{(a)}} - F_d(Y) \right) \left(\frac{1}{\overline{\rho(Y|\cdot)^{(a)}}} - \frac{1}{\rho(Y)} \right) \\
 &+ \frac{1}{2} \left(\overline{f_d(\cdot)\rho(Y|\cdot)^{(b)}} - F_d(Y) \right) \left(\frac{1}{\overline{\rho(Y|\cdot)^{(b)}}} - \frac{1}{\rho(Y)} \right) \\
 &- \left(\overline{f_d(\cdot)\rho(Y|\cdot)} - F_d(Y) \right) \left(\frac{1}{\overline{\rho(Y|\cdot)}} - \frac{1}{\rho(Y)} \right) \\
 &+ \frac{1}{2} \frac{F_d(Y)}{\rho(Y)} \left(\frac{\rho(Y)}{\overline{\rho(Y|\cdot)^{(a)}}} - 2 + \frac{\overline{\rho(Y|\cdot)^{(a)}}}{\rho(Y)} \right) \\
 &+ \frac{1}{2} \frac{F_d(Y)}{\rho(Y)} \left(\frac{\rho(Y)}{\overline{\rho(Y|\cdot)^{(b)}}} - 2 + \frac{\overline{\rho(Y|\cdot)^{(b)}}}{\rho(Y)} \right) \\
 &- \frac{F_d(Y)}{\rho(Y)} \left(\frac{\rho(Y)}{\overline{\rho(Y|\cdot)}} - 2 + \frac{\overline{\rho(Y|\cdot)}}{\rho(Y)} \right).
 \end{aligned}$$

If $|\overline{\rho(Y|\cdot)^{(a)}}/\rho(Y) - 1| < 1/2$ and $|\overline{\rho(Y|\cdot)^{(b)}}/\rho(Y) - 1| < 1/2$, it follows that $|\overline{\rho(Y|\cdot)}/\rho(Y) - 1| < 1/2$, and therefore by applying Jensen's inequality we obtain

$$\begin{aligned}
 \left| \frac{1}{2} \left(\overline{g_d^{(a)}} + \overline{g_d^{(b)}} \right) - \overline{g_d} \right|^q &\leq 2^{2q-3} \left| \overline{f_d(\cdot)\rho(Y|\cdot)^{(a)}} - F_d(Y) \right|^q \left| \frac{1}{\overline{\rho(Y|\cdot)^{(a)}}} - \frac{1}{\rho(Y)} \right|^q \\
 &+ 2^{2q-3} \left| \overline{f_d(\cdot)\rho(Y|\cdot)^{(b)}} - F_d(Y) \right|^q \left| \frac{1}{\overline{\rho(Y|\cdot)^{(b)}}} - \frac{1}{\rho(Y)} \right|^q \\
 &+ 2^{2q-2} \left| \overline{f_d(\cdot)\rho(Y|\cdot)} - F_d(Y) \right|^q \left| \frac{1}{\overline{\rho(Y|\cdot)}} - \frac{1}{\rho(Y)} \right|^q \\
 &+ 2^{2q-3} \frac{|F_d(Y)|^q}{\rho(Y)^q} \left| \frac{\rho(Y)}{\overline{\rho(Y|\cdot)^{(a)}}} - 2 + \frac{\overline{\rho(Y|\cdot)^{(a)}}}{\rho(Y)} \right|^q \\
 &+ 2^{2q-3} \frac{|F_d(Y)|^q}{\rho(Y)^q} \left| \frac{\rho(Y)}{\overline{\rho(Y|\cdot)^{(b)}}} - 2 + \frac{\overline{\rho(Y|\cdot)^{(b)}}}{\rho(Y)} \right|^q
 \end{aligned}$$

$$\begin{aligned}
& + 2^{2q-2} \frac{|F_d(Y)|^q}{\rho(Y)^q} \left| \frac{\rho(Y)}{\rho(Y|\cdot)} - 2 + \frac{\overline{\rho(Y|\cdot)}}{\rho(Y)} \right|^q \\
& \leq 2^{3q-3} \left| \overline{f_d(\cdot)\rho(Y|\cdot)}^{(a)} - F_d(Y) \right|^q \left| \frac{\overline{\rho(Y|\cdot)}^{(a)}}{\rho(Y)} - 1 \right|^q \frac{1}{\rho(Y)^q} \\
& \quad + 2^{3q-3} \left| \overline{f_d(\cdot)\rho(Y|\cdot)}^{(b)} - F_d(Y) \right|^q \left| \frac{\overline{\rho(Y|\cdot)}^{(b)}}{\rho(Y)} - 1 \right|^q \frac{1}{\rho(Y)^q} \\
& \quad + 2^{3q-2} \left| \overline{f_d(\cdot)\rho(Y|\cdot)} - F_d(Y) \right|^q \left| \frac{\overline{\rho(Y|\cdot)}}{\rho(Y)} - 1 \right|^q \frac{1}{\rho(Y)^q} \\
& \quad + 2^{3q-3} |F_d(Y)|^q \left| \frac{\overline{\rho(Y|\cdot)}^{(a)}}{\rho(Y)} - 1 \right|^{2q} \frac{1}{\rho(Y)^q} \\
& \quad + 2^{3q-3} |F_d(Y)|^q \left| \frac{\overline{\rho(Y|\cdot)}^{(b)}}{\rho(Y)} - 1 \right|^{2q} \frac{1}{\rho(Y)^q} \\
& \quad + 2^{3q-2} |F_d(Y)|^q \left| \frac{\overline{\rho(Y|\cdot)}}{\rho(Y)} - 1 \right|^{2q} \frac{1}{\rho(Y)^q}.
\end{aligned}$$

Looking at the third term, if $p \geq 2q$, then Hölder's inequality, together with Lemma 3.8, (3.3), and (3.4), gives

$$\begin{aligned}
& \mathbb{E} \left[\left| \overline{f_d(\cdot)\rho(Y|\cdot)} - F_d(Y) \right|^q \left| \frac{\overline{\rho(Y|\cdot)}}{\rho(Y)} - 1 \right|^q \frac{1}{\rho(Y)^q} \right] \\
& \leq \left(\mathbb{E} \left[\left| \overline{f_d(\cdot)\rho(Y|\cdot)} - F_d(Y) \right|^{2q} / \rho(Y)^{2q} \right] \mathbb{E} \left[\left| \frac{\overline{\rho(Y|\cdot)}}{\rho(Y)} - 1 \right|^{2q} \right] \right)^{1/2} \\
& \leq 2^{q-1/2} C_{2q} 2^{-q\ell} M_0^{-q} F_{max}^q \left(\mathbb{E}_Y \left[\mathbb{E}_\theta \left[\left(\frac{\rho(Y|\theta)}{\rho(Y)} \right)^{2q} \right] \right] + 1 \right).
\end{aligned}$$

If $2 \leq p < 2q$, then we need to modify the argument slightly, using Hölder's inequality when $q < p$ and a simple identity when $q = p$, to obtain

$$\begin{aligned}
& \mathbb{E} \left[\left| \overline{f_d(\cdot)\rho(Y|\cdot)} - F_d(Y) \right|^q \left| \frac{\overline{\rho(Y|\cdot)}}{\rho(Y)} - 1 \right|^q \frac{1}{\rho(Y)^q} \mathbf{1}_{E^c} \right] \\
& \leq 2^{p-2q} \mathbb{E} \left[\left| \overline{f_d(\cdot)\rho(Y|\cdot)} - F_d(Y) \right|^q \left| \frac{\overline{\rho(Y|\cdot)}}{\rho(Y)} - 1 \right|^{p-q} \frac{1}{\rho(Y)^q} \right] \\
& \leq 2^{p-2q} \left(\mathbb{E} \left[\left| \overline{f_d(\cdot)\rho(Y|\cdot)} - F_d(Y) \right|^p / \rho(Y)^p \right] \right)^{q/p} \left(\mathbb{E} \left[\left| \frac{\overline{\rho(Y|\cdot)}}{\rho(Y)} - 1 \right|^p \right] \right)^{1-q/p} \\
& \leq 2^{p-q-q/p} C_p 2^{-p\ell/2} M_0^{-p/2} F_{max}^q \left(\mathbb{E}_Y \left[\mathbb{E}_\theta \left[\left(\frac{\rho(Y|\theta)}{\rho(Y)} \right)^p \right] \right] + 1 \right).
\end{aligned}$$

Looking at the sixth term, by Lemma 3.8 and (3.3) we obtain

$$\begin{aligned} & \mathbb{E} \left[|F_d(Y)|^q \left| \frac{\rho(Y|\cdot)}{\rho(Y)} - 1 \right|^{2q} \frac{1}{\rho(Y)^q} \right] \\ & \leq C_{2q} 2^{-q\ell} M_0^{-q} F_{max}^q \left(\mathbb{E}_Y \left[\mathbb{E}_\theta \left[\left(\frac{\rho(Y|\theta)}{\rho(Y)} \right)^{2q} \right] \right] + 1 \right), \end{aligned}$$

when $p \geq 2q$, while for $2 \leq p < 2q$ we get

$$\begin{aligned} & \mathbb{E} \left[|F_d(Y)|^q \left| \frac{\rho(Y|\cdot)}{\rho(Y)} - 1 \right|^{2q} \frac{1}{\rho(Y)^q} \mathbf{1}_{E^c} \right] \\ & \leq 2^{p-2q} C_p 2^{-p\ell/2} M_0^{-p/2} F_{max}^q \left(\mathbb{E}_Y \left[\mathbb{E}_\theta \left[\left(\frac{\rho(Y|\theta)}{\rho(Y)} \right)^p \right] \right] + 1 \right). \end{aligned}$$

The other terms can be treated similarly, and therefore, adding in the bound from (3.5), we determine that there exists a constant C , dependent only on p and q , such that

$$\begin{aligned} & \mathbb{E} \left[\left| \frac{1}{2} \left(\bar{g}_d^{(a)} + \bar{g}_d^{(b)} \right) - \bar{g}_d \right|^q \right] \\ & \leq C 2^{-\min(p,2q)\ell/2} M_0^{-\min(p,2q)/2} F_{max}^q \left(\mathbb{E}_Y \left[\mathbb{E}_\theta \left[\left(\frac{\rho(Y|\theta)}{\rho(Y)} \right)^{\min(p,2q)} \right] \right] + 1 \right), \end{aligned}$$

completing the proof of the second assertion. ■

Remark 3.10. It follows from Assumption 3.1 that

$$\left| \max_{d \in D} f_d(\theta) \right| \leq F_{max}$$

for any θ . Given that the antithetic property (2.4) holds, the same proof strategy leads to

$$\mathbb{E}[|\Delta P_{\ell,1}|^q] = \mathbb{E} \left[\left| \bar{g}_{\max} - \frac{1}{2} \left(\bar{g}_{\max}^{(a)} + \bar{g}_{\max}^{(b)} \right) \right|^q \right] = O(2^{-\min(p,2q)\ell/2}),$$

where p is as defined in Assumption 3.4.

Remark 3.11. Without the antithetic technique, the best that can be deduced from Lemma 3.9 is that

$$\begin{aligned} \mathbb{E} \left[|\bar{g}_d^{(a)} - \bar{g}_d|^q \right] &= \mathbb{E} \left[|\bar{g}_d^{(a)} - G_d - \bar{g}_d + G_d|^q \right] \\ &\leq 2^{q-1} \mathbb{E} \left[|\bar{g}_d^{(a)} - G_d|^q \right] + 2^{q-1} \mathbb{E} \left[|\bar{g}_d - G_d|^q \right] = O(2^{-q\ell/2}) \end{aligned}$$

for $1 \leq q \leq p$. This leads to the variance of the nonantithetic correction $\Delta P_\ell = P_\ell - P_{\ell-1}^{(a)}$ being $O(2^{-\ell})$. Since we now have $\beta = 1$ in Theorem 2.1, the nonantithetic MLMC estimator is in the second regime regardless of the value of $p \geq 2$. The resulting total computational complexity is thus $O(\varepsilon^{-2}(\log \varepsilon^{-1})^2)$, which is not optimal.

4. Numerical experiments.

4.1. Problem setting. To demonstrate the efficiency of our MLMC estimator, we conduct some numerical experiments using a modified version of the model originally introduced in [1] and explored further in [25]. The original version of this artificial cost-effectiveness model compares a new treatment to some standard of care on the prevention of a critical event, denoted by E . We extend this to three decision options, $D = \{1, 2, 3\}$. We also modify the model to include correlations between odds ratios (OR) of the critical events ($OR_{E,2}$, $OR_{E,3}$), treatment costs ($C_{T,2}$, $C_{T,3}$), and probabilities of side effects ($P_{SE,2}$, $P_{SE,3}$). The resulting vector θ consists of 12 random variables, L , Q_E , Q_{SE} , C_E , C_{SE} , $C_{T,2}$, $C_{T,3}$, $P_{E,1}$, $OR_{E,2}$, $OR_{E,3}$, $P_{SE,2}$, $P_{SE,3}$. The model contains only three constants, $C_{T,1}$, $P_{SE,1}$, λ , while two parameters, $P_{E,2}$ and $P_{E,3}$, are defined as functions of $P_{E,1}$, $OR_{E,2}$ and $P_{E,1}$, $OR_{E,3}$, respectively. We refer the reader to Table 1 for a detailed description of these model inputs. For each treatment $d \in D$, the net benefit is defined by

$$\begin{aligned} f_d(\theta) = & P_{SE,d}P_{E,d} \left[\lambda \left(L \frac{1+Q_E}{2} - Q_{SE} \right) - (C_{SE} + C_E) \right] \\ & + P_{SE,d}(1 - P_{E,d}) [\lambda(L - Q_{SE}) - C_{SE}] \\ & + (1 - P_{SE,d})P_{E,d} \left[\lambda L \frac{1+Q_E}{2} - C_E \right] \\ & + (1 - P_{SE,d})(1 - P_{E,d})\lambda L - C_{T,d}. \end{aligned}$$

This net benefit, as in [1], is multilinear, with the consequence that expected values of random variables can be plugged in to form estimates of the EVSI. However, the additional decision options, correlations, and larger number of parameters make the problem more realistic, and computationally challenging, than the original version. Repeating the standard Monte Carlo estimation with 10^8 i.i.d. samples of θ 10 times, EVPI is estimated as 4,063.5 with the standard error equal to 0.66.

Our final modification to improve practical relevance is to assume an annual population that experiences this disease as 2,500, a technology horizon of 10 years, and an annual discount factor of 1.035, giving a total discounted population that will benefit from sampling as 21,519. Estimated per-person EVSI is multiplied by this factor to get a population EVSI which can be compared with experiment costs.

We consider three hypothetical experiment designs for which we want to estimate the EVSI. Unlike [1, 25], all experiments inform more than one parameter, and we consider cases where the parameters being studied (subset of θ informed by sample information Y) are correlated with those not being studied (subset of θ not informed by Y). These modifications again improve the practical relevance of the example.

1. An observational study on side effects in a cohort of $n_p = 100$ patients on treatment $d = 2$, which informs three parameters Q_{SE} , C_{SE} , and $P_{SE,2}$. The sample information Y in this scenario, denoted by Scenario 1, is defined as the three-dimensional vector (Y_1, Y_2, Y_3) , where $Y_1 \sim N(Q_{SE}, 4/n_p)$, $Y_2 \sim N(C_{SE}, 10^4/n_p)$, and $Y_3 \sim \text{Binomial}(n_p, P_{SE,2})$. The hypothetical study costs are \$50,000 for setup and \$250 for each patient giving a total of \$75,000.

Table 1

Model inputs involved in our experiments. Note that $\log\text{-normal}(\mu, \Sigma)$ and $\text{logit-normal}(\mu, \Sigma)$ denote the log-normal and logit-normal distributions, respectively, with μ and Σ being the mean vector and the covariance matrix of the corresponding normal distribution, respectively. The logit function $\text{logit}(p)$ is the transformation mapping probability p to the logarithm of the odds $\log(p/(1-p))$.

Description	Parameter	Distribution
Lifetime remaining	L	$N(30, 25)$
QALY after critical event, per year	Q_E	$\text{logit-normal}(0.6, 1/36)$
QALY decrement due to side effects	Q_{SE}	$N(0.7, 0.01)$
Cost of critical event	C_E	$N(2 \times 10^5, 10^8)$
Cost of side effect	C_{SE}	$N(10^5, 10^8)$
Cost of treatment $d = 1$	$C_{T,1}$	0 (constant)
Cost of treatments $d = 2, 3$	$C_{T,d}$	$N\left(\begin{pmatrix} 1.5 \times 10^4 \\ 2 \times 10^4 \end{pmatrix}, \begin{pmatrix} 300 & 100 \\ 100 & 500 \end{pmatrix}\right)$
Probability of critical event on treatment $d = 1$	$P_{E,1}$	Beta(15, 85)
Odds ratios of critical event relative to treatment $d = 1$ ($P_{E,d}/(1 - P_{E,d})/(P_{E,1}/(1 - P_{E,1}))$)	$OR_{E,d}$	$\log\text{-normal}\left(\begin{pmatrix} -1.5 \\ -1.75 \end{pmatrix}, \begin{pmatrix} 0.11 & 0.02 \\ 0.02 & 0.06 \end{pmatrix}\right)$
Probability of critical event on treatments $d = 2, 3$	$P_{E,d}$	Derived from $P_{E,1}$ and $OR_{E,d}$
Probability of side effect on treatment $d = 1$	$P_{SE,1}$	0 (constant)
Probability of side effect on treatments $d = 2, 3$	$P_{SE,d}$	$\text{logit-normal}\left(\begin{pmatrix} -1.4 \\ -1.1 \end{pmatrix}, \begin{pmatrix} 0.10 & 0.05 \\ 0.05 & 0.25 \end{pmatrix}\right)$
Monetary value of 1 QALY	λ	\$75,000 (constant)

2. A small two-arm randomized controlled trial (RCT) comparing treatments $d = 1$ and $d = 3$ with $n_p = 100$ patients, which informs three parameters $OR_{E,3}$, $C_{T,3}$, and $P_{SE,3}$. The sample information Y in this scenario, denoted by Scenario 2, is defined as the three-dimensional vector (Y_1, Y_2, Y_3) , where $Y_1 \sim N(\log(OR_{E,3}), 4/n_p)$, $Y_2 \sim N(C_{T,3}, 10^4/n_p)$, and $Y_3 \sim \text{Binomial}(n_p, P_{SE,3})$. This hypothetical RCT costs \$200,000 for setup plus \$1000 for each randomized patient, giving a total of \$400,000.
3. A larger multicenter two-arm randomized controlled trial comparing treatments $d = 1$ and $d = 3$ with $n_p = 1000$ patients, which informs six parameters $P_{E,1}$, $OR_{E,3}$, $C_{T,3}$, $P_{SE,3}$, C_{SE} , and C_E . The sample information Y in this scenario, denoted by Scenario 3, is defined as the six-dimensional vector (Y_1, \dots, Y_6) , where $Y_1 \sim \text{Binomial}(n_p, P_{E,1})$, $Y_2 \sim N(\log(OR_{E,3}), 4/n_p)$, $Y_3 \sim N(C_{T,3}, 10^4/n_p)$, $Y_4 \sim \text{Binomial}(n_p, P_{SE,3})$, $Y_5 \sim N(C_{SE}, 10^4/n_p)$, and $Y_6 \sim N(C_E, 10^4/n_p)$. This hypothetical multicenter RCT costs \$200,000 for setup, \$1500 for each randomized patient, and \$50,000 for each additional 100 patients added to the study (owing to the cost of establishing a new center). The total cost is \$4,200,000.

4.2. Results and discussion. Throughout our experiments below, we always set $M_0 = 16$. Here we refer the reader to [12, p. 516] for the discussion on how to choose M_0 to minimize the total complexity of the MLMC estimator. Except for $P_{SE,2}$ for Scenario 1 and $P_{SE,3}$ for Scenarios 2 and 3, the exact posterior distributions of the relevant random input parameters

given Y are available and used as importance distributions. Regarding $P_{SE,2}$ for Scenario 1, we approximate the prior distribution by the beta distribution with the same mean and variance, for which the exact posterior distribution given Y_3 is available and can be used as an importance distribution. We employ the same approach for $P_{SE,3}$ for Scenarios 2 and 3.

Following the MLMC implementation by Giles [9, section 3], we first conduct convergence tests on a sequence of the correction variables ΔP_ℓ and then estimate EVPI – EVSI for various values of the root-mean-square accuracy ε . The required computational costs for our MLMC estimator are compared with those for the standard nested Monte Carlo estimator of the form

$$\frac{1}{N} \sum_{n=1}^N P_L^{(n)}$$

for the same maximum level L with the MLMC estimator. Here we note that this nested Monte Carlo estimator is different from that given in (1.5). We also note that function approximation techniques [25, 14, 15] may be applied to this example. However, performance comparison with the nested MCMC estimator or function approximation-based estimators is beyond the scope of this article and will be addressed in future work.

The results for Scenario 1 are shown in Figure 1. In the left top panel, the variances of both P_ℓ and ΔP_ℓ are plotted as functions of the level ℓ . Here the variances are estimated empirically by using 10^5 i.i.d. samples at each level. While the variance of P_ℓ takes an almost constant value, the variance of ΔP_ℓ decreases geometrically as the level increases. The linear regression of the data for the range $2 \leq \ell \leq 8$ provides an estimation of β as 1.62, which is slightly better than what we expect from the theoretical analysis. Using the same i.i.d. samples of P_ℓ and ΔP_ℓ , we also estimate their absolute mean values, which are plotted as functions of ℓ in the right top panel. The absolute mean value of P_ℓ takes an almost constant value, whereas the absolute mean value of ΔP_ℓ decreases geometrically as the level increases. The linear regression of the data for the range $2 \leq \ell \leq 8$ provides an estimation of α as 1.16, which again is slightly better than our theoretical result. These convergence results confirm that our MLMC estimator is in the first regime of Theorem 2.1 for this scenario.

For a given root-mean-square accuracy ε , the necessary maximum level L and the corresponding numbers of samples N_ℓ for $\ell = 0, \dots, L$ are estimated by using [9, Algorithm 1]. Each line in the left bottom panel of Figure 1 shows the values of N_ℓ for a particular value of $\varepsilon \in \{0.5, 1, 2, 5, 10, 20\}$. It can be seen that the required maximum level L increases as ε decreases so that the bias becomes small enough. For each fixed ε , the number of samples N_ℓ decreases geometrically with the level, and most of the samples are allocated on the lower levels. In fact, the optimal allocation of N_ℓ is given by $N_\ell \propto \varepsilon^{-2} 2^{-(\beta+\gamma)\ell/2}$ [9], and our experimental results agree quite well with this.

It should be noted that the values of both α and β are estimated on the fly as the computation is performed, with the value of α being used to determine when the bias $\mathbb{E}[P - P_L]$ has converged sufficiently as L increases. Further details on the MLMC software implementation are available in [9].

The total computational cost for the MLMC estimator is given by the sum $\sum_{\ell=0}^L 2^\ell N_\ell =: C_{\text{MLMC}}$, and it is expected from Theorem 2.1 and the above convergence results that $\varepsilon^2 C_{\text{MLMC}}$ is independent of ε . On the other hand, the total computational cost for the nested Monte

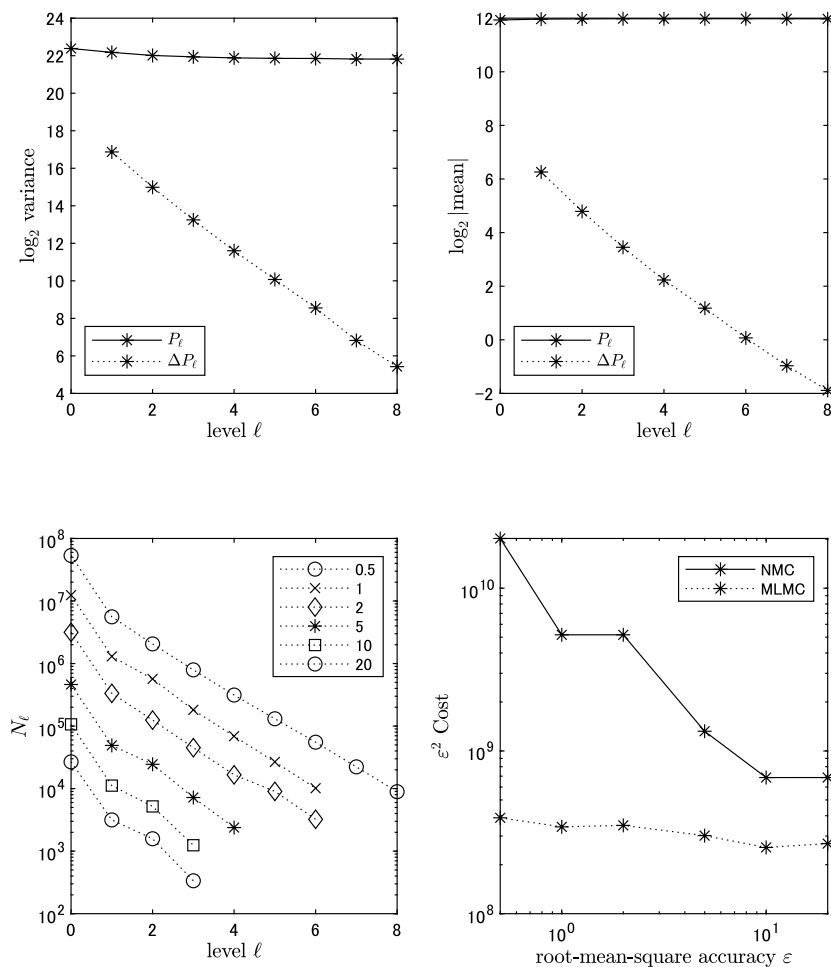


Figure 1. MLMC results for Scenario 1.

Carlo estimator, denoted by C_{NMC} , is of order $\epsilon^{-2-1/\alpha}$ theoretically, so that $\epsilon^2 C_{\text{NMC}}$ should be proportional to $\epsilon^{-1/\alpha}$ and depends strongly on ϵ . Based on this observation, the values of $\epsilon^2 C_{\text{MLMC}}$ and $\epsilon^2 C_{\text{NMC}}$ are plotted as functions of ϵ with dashed and solid lines, respectively, in the right bottom panel. As expected, $\epsilon^2 C_{\text{NMC}}$ increases significantly as ϵ decreases, while $\epsilon^2 C_{\text{MLMC}}$ takes an almost constant value. As a consequence, the superiority of the MLMC estimator becomes more remarkable for smaller ϵ . For the smallest value $\epsilon = 0.5$ in this experiment, the MLMC estimator achieves computational saving of factor 52. This factor is expected to get even larger if we set a smaller value for ϵ .

We obtain similar results for Scenarios 2 and 3 as shown in Figures 2 and 3, respectively. The estimated values of α and β are 0.99 and 1.48, respectively, for both these scenarios, which agrees quite well with our theoretical finding. The MLMC estimator achieves computational saving of factors 35 and 118 for the smallest value $\epsilon = 0.5$. The estimates of the difference $\text{EVPI} - \text{EVSI}$ are 4,039, 3,033, and 2,277 for the three scenarios, respectively, with the root-mean-square accuracy equal to 0.5. Subtracting these values from the estimate of EVPI, which

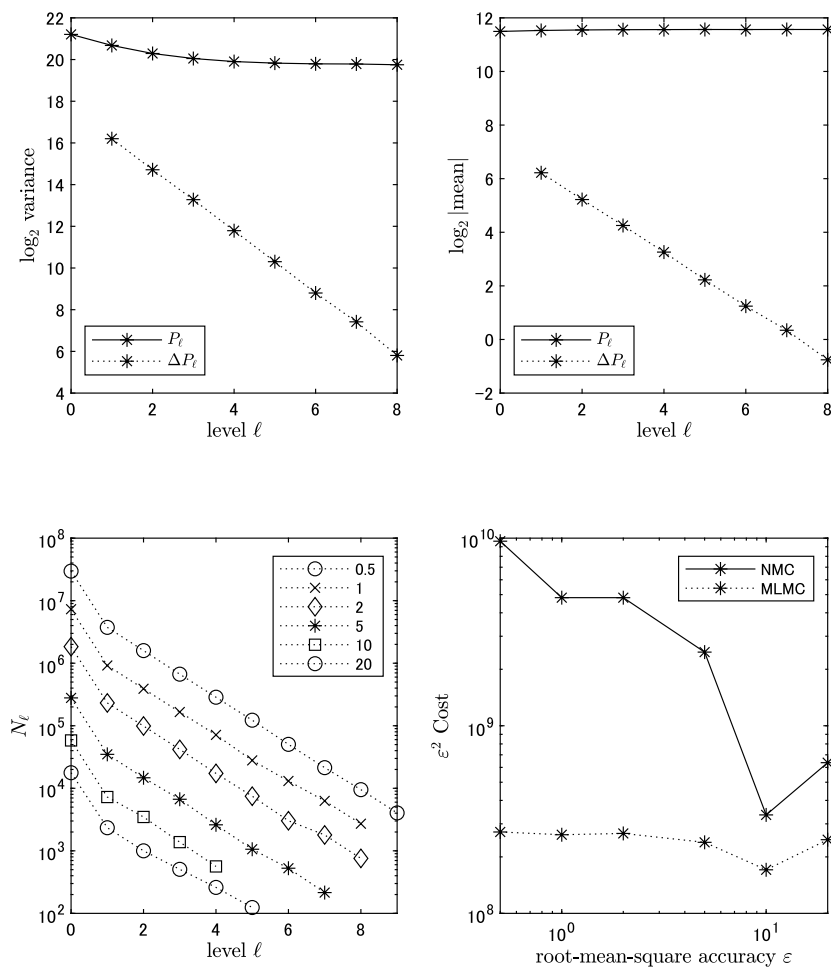


Figure 2. MLMC results for Scenario 2.

is 4,064, the per-person EVSI is estimated as 25, 1,031, and 1,787, respectively. Although we round the estimates to the nearest integers here for simplicity, each EVSI is estimated with the root-mean-square accuracy at most $0.66 + 0.5 = 1.06$. As comparison, by using the MLMC estimator for EVPPI [11], the per-person EVPPI for $(Q_{SE}, C_{SE}, P_{SE,2})$, $(OR_{E,3}, C_{T,3}, P_{SE,3})$, and $(P_{SE,1}, OR_{E,3}, C_{T,3}, P_{SE,3}, C_{SE}, C_E)$ is estimated as 86, 1,400, and 1,914, each of which is certainly larger than the corresponding EVSI, confirming face validity of our EVSI estimates.

Multiplying the EVSI estimates by the discounted population that will benefit from sampling gives the population EVSI of \$537,975, \$22,186,089, and \$38,454,453 (with the root-mean-square accuracy at most $1.06 \times 21,519 = 22,810$) for Scenarios 1, 2, and 3, respectively, which are all larger than their hypothetical costs and are thus cost-effective. Subtracting the study costs gives the expected net benefit of sampling (ENBS) [27]. In our three scenarios, the ENBS values are estimated as \$462,975, \$21,786,089, and \$34,254,453, respectively. The ENBS for the third experiment, the large multicenter randomized controlled trial, is the highest and is thus the study that should be funded.

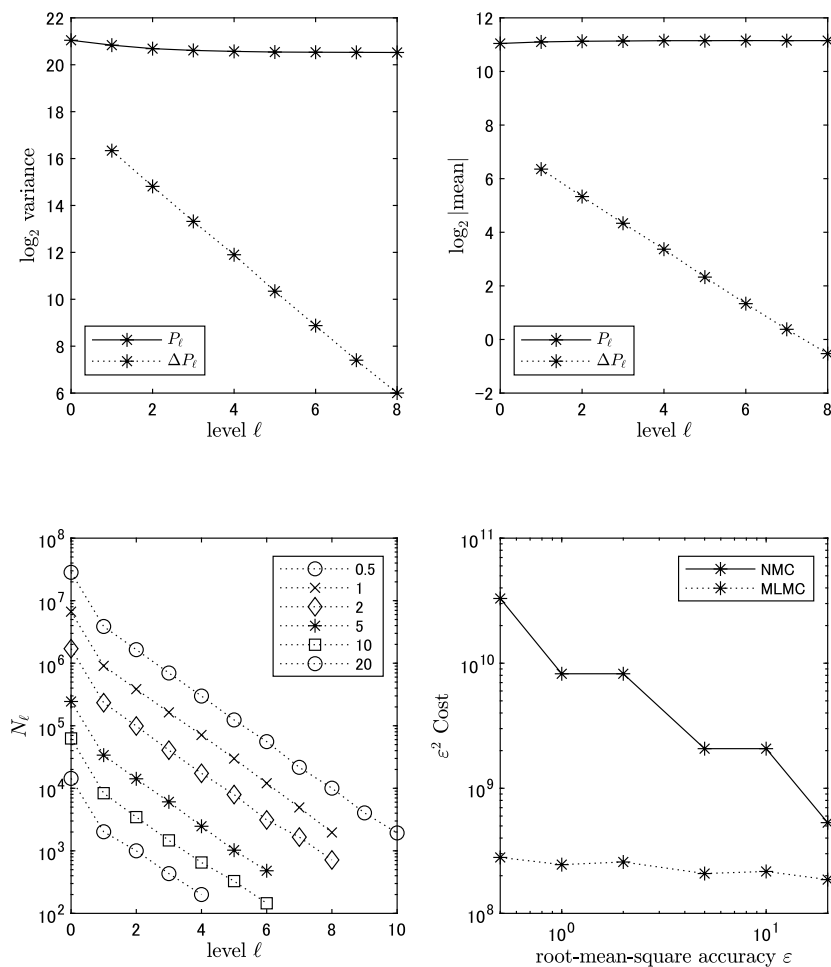


Figure 3. MLMC results for Scenario 3.

5. Conclusions. In this paper, we have developed a multilevel Monte Carlo estimator for EVSI, the expected value of sample information. The key difference from the multilevel estimation for EVPPI, as studied in [11], is to use Bayes' theorem directly to rewrite the expectation with respect to the posterior distribution of input random variables θ given an observation Y into the ratio of the expectations with respect to the prior distribution of θ , and then to estimate each of the expectations by using the same i.i.d. samples of θ . As shown in Lemma 3.9, we prove under Assumptions 3.1 and 3.4 that the nested ratio expectation can be efficiently estimated by using the antithetic multilevel Monte Carlo estimator. Plugging this result into a slightly generalized version of [11, Theorem 3], our antithetic multilevel estimator is shown to achieve a root-mean-square accuracy ε at a cost of optimal $O(\varepsilon^{-2})$. As mentioned in Remark 3.11, without the antithetic technique, we can only expect a suboptimal result for the overall computational complexity. Our theoretical analysis is supported by numerical experiments.

In future work, following the idea from [3, 12], we will examine the use of an adaptive

number of inner samples for θ with an aim to bring about further computational savings for estimating EVSI. For the outer samples of Y which are far away from the decision manifold K , a smaller number of inner samples than $M_0 2^\ell$ may be sufficient to ensure that

$$\arg \max_{d \in D} \bar{g}_d = \arg \max_{d \in D} \bar{g}_d^{(a)} = \arg \max_{d \in D} \bar{g}_d^{(b)} = d_{\text{opt}}$$

holds with high probability. Thus it is only the outer samples of Y which are close to K that require great accuracy for estimating the inner conditional expectation. In this way we expect to reduce the value of γ while maintaining the fast decay of $\mathbb{E}[|\Delta P_\ell|]$ and $\mathbb{V}[\Delta P_\ell]$ as derived in Theorem 3.6. This should reduce the total computational cost by some constant factor, approximately independent of the desired accuracy.

Acknowledgment. The authors would like to thank Abdul-Lateef Haji-Ali of Heriot-Watt University for useful discussions and comments on the earlier version of this article.

REFERENCES

- [1] A. E. ADES, G. LU, AND K. CLAXTON, *Expected value of sample information calculations in medical decision modeling*, Medical Decision Making, 24 (2004), pp. 207–227.
- [2] A. BRENNAN, S. KHARROUBI, A. O’HAGAN, AND J. CHILCOTT, *Calculating partial expected value of perfect information via Monte Carlo sampling algorithms*, Medical Decision Making, 27 (2007), pp. 448–470.
- [3] M. BROADIE, Y. DU, AND C. MOALLEMI, *Efficient risk estimation via nested sequential simulation*, Management Sci., 57 (2011), pp. 1172–1194.
- [4] K. BUJOK, B. HAMBLY, AND C. REISINGER, *Multilevel simulation of functionals of Bernoulli random variables with application to basket credit derivatives*, Methodol. Comput. Appl. Probab., 17 (2015), pp. 579–604.
- [5] K. A. CLIFFE, M. B. GILES, R. SCHEICHL, AND A. L. TECKENTRUP, *Multilevel Monte Carlo methods and applications to elliptic PDEs with random coefficients*, Comput. Vis. Sci., 14 (2011), pp. 3–15.
- [6] J. DICK, R. N. GANTNER, Q. T. L. GIA, AND C. SCHWAB, *Multilevel higher-order quasi-Monte Carlo Bayesian estimation*, Math. Models Methods Appl. Sci., 27 (2018), pp. 953–995.
- [7] R. N. GANTNER AND M. D. PETERS, *Higher-order quasi-Monte Carlo for Bayesian shape inversion*, SIAM/ASA J. Uncertainty Quantif., 6 (2018), pp. 707–736, <https://doi.org/10.1137/16M1096116>.
- [8] M. B. GILES, *Multilevel Monte Carlo path simulation*, Oper. Res., 56 (2008), pp. 607–617.
- [9] M. B. GILES, *Multilevel Monte Carlo methods*, Acta Numer., 24 (2015), pp. 259–328.
- [10] M. B. GILES, *MLMC for nested expectations*, in Contemporary Computational Mathematics: A Celebration of the 80th Birthday of Ian Sloan, Springer, 2018, pp. 425–442.
- [11] M. B. GILES AND T. GODA, *Decision-making under uncertainty: Using MLMC for efficient estimation of EVPPI*, Stat. Comput., 29 (2019), pp. 739–751.
- [12] M. B. GILES AND A.-L. HAJI-ALI, *Multilevel nested simulation for efficient risk estimation*, SIAM/ASA J. Uncertainty Quantif., 7 (2019), pp. 497–525, <https://doi.org/10.1137/18M1173186>.
- [13] T. GODA, T. HIRONAKA, AND T. IWAMOTO, *Multilevel Monte Carlo estimation of expected information gains*, Stoch. Anal. Appl., 38 (2020), pp. 581–600.
- [14] A. HEATH, I. MANOLOPOULOU, AND G. BAIO, *Estimating the expected value of partial perfect information in health economic evaluations using integrated nested Laplace approximation*, Statist. Medicine, 35 (2016), pp. 4264–4280.
- [15] H. JALAL AND F. ALARID-ESCUADERO, *A Gaussian approximation approach for value of information analysis*, Medical Decision Making, 38 (2018), pp. 174–188.
- [16] A. JOULIN AND Y. OLLIVIER, *Curvature, concentration and error estimates for Markov chain Monte Carlo*, Ann. Probab., 38 (2010), pp. 2418–2442.
- [17] K. ŁATUSZYŃSKI, B. MIASOJEDOW, AND W. NIEMIRO, *Nonasymptotic bounds on the estimation error of MCMC algorithms*, Bernoulli, 19 (2013), pp. 2033–2066.

- [18] J. S. LIU, *Monte Carlo Strategies in Scientific Computing*, Springer, 2004.
- [19] N. A. MENZIES, *An efficient estimator for the expected value of sample information*, *Medical Decision Making*, 36 (2016), pp. 308–320.
- [20] C. P. ROBERT AND G. CASELLA, *Monte Carlo Statistical Methods*, 2nd ed., Springer, 2004.
- [21] R. SCHEICHL, A. M. STUART, AND A. L. TECKENTRUP, *Quasi-Monte Carlo and multilevel Monte Carlo methods for computing posterior expectations in elliptic inverse problems*, *SIAM/ASA J. Uncertainty Quantif.*, 5 (2017), pp. 493–518, <https://doi.org/10.1137/16M1061692>.
- [22] C. SCHILLINGS AND C. SCHWAB, *Sparse, adaptive Smolyak quadratures for Bayesian inverse problems*, *Inverse Problems*, 29 (2013), pp. 1–28.
- [23] C. SCHILLINGS AND C. SCHWAB, *Scaling limits in computational Bayesian inversion*, *ESAIM Math. Model. Numer. Anal.*, 50 (2016), pp. 1825–1856.
- [24] C. SCHILLINGS, B. SPRUNGK, AND P. WACKER, *On the convergence of the Laplace approximation and noise-level-robustness of Laplace-based Monte Carlo methods for Bayesian inverse problems*, *Numer. Math.*, 145 (2020), pp. 915–971.
- [25] M. STRONG, J. E. OAKLEY, A. BRENNAN, AND P. BREEZE, *Estimating the expected value of sample information using the probabilistic sensitivity analysis sample: A fast, nonparametric regression-based method*, *Medical Decision Making*, 35 (2015), pp. 570–583.
- [26] A. M. STUART, *Inverse problems: A Bayesian perspective*, *Acta Numer.*, 19 (2010), pp. 451–559.
- [27] N. J. WELTON, A. J. SUTTON, N. J. COPPER, K. R. ABRAMS, AND A. E. ADES, *Evidence Synthesis for Decision Making in Healthcare*, Wiley, 2012.