NO SIGMOID REQUIRED: EXACT FEASIBILITY FOR CONSTRAINED qEHVI

Anonymous authors
Paper under double-blind review

ABSTRACT

Constrained multi-objective optimization is essential for engineering design, requiring sample-efficient batch evaluation strategies due to costly simulations. Although Bayesian optimization using qCEHVI is highly effective, estimating its probability of feasibility (PoF) relies on a sigmoid relaxation, which introduces bias and a tunable parameter. Naïve marginalization can eliminate the sigmoid, but we find that it can lead to overestimation of the acquisition function for $Q > 1$ when running points are spatially concentrated. This paper theoretically analyzes this error's source and derives an unbiased, low-variance Rao-Blackwellized (RB) estimator that replaces the candidate feasibility indicator with its exact conditional expectation. Using a synthetic problem that reproduces a concentrated scenario, we first verify that the RB approach suppresses the overestimation in the infeasible region. Furthermore, our approach resolves the sigmoid-induced PoF overestimation, improving sample efficiency near the feasibility boundary. Experiments on real-world benchmark problems demonstrate that our method yields stable performance, eliminating the parameter sensitivity of the sigmoid relaxation.

1 INTRODUCTION

Constrained multi-objective optimization (CMO) is pervasive in engineering design configurations. It arises whenever a configuration must satisfy various physical, performance, and manufacturing constraints while balancing competing objectives like safety versus efficiency and cost versus performance. Examples include structural design (Coello Coello & Pulido, 2005; Ray & Liew, 2002), mechanical design (Farhang-Mehr & Azarm, 2002; Deb & Srinivasan, 2006), automotive and marine engineering (Jain & Deb, 2013; Parsons & Scott, 2004), and civil infrastructure planning (Ray et al., 2001). Since evaluating these configurations often requires costly simulations or physical tests, candidate configurations are typically evaluated in parallel as a batch. This motivates the need for an efficient batch suggestion strategy that quickly identifies feasible Pareto-optimal configurations.

While CMO has historically been led by the evolutionary algorithm community (Coello, 2002; Liang et al., 2022; Rahimi et al., 2023), recent studies claim that Bayesian optimization (BO) is more sample-efficient than evolutionary algorithms (Shahriari et al., 2015; Daulton et al., 2022; Renganathan & Carlson, 2025), driving the emerging adoption of BO in CMO (Hernández-Lobato et al., 2015; Daulton et al., 2020; Ament et al., 2023). BO generally quantifies the promise of each configuration choice using an acquisition function (AF). Some of the most widely adopted AFs belong to the expected improvement family (Ament et al., 2023), such as constrained expected hypervolume improvement (CEHVI) for CMO, which computes the product of the probability of feasibility (PoF) and the expected hypervolume improvement (EHVI) (Daulton et al., 2020). CEHVI can be extended to the batch suggestion scenario, where a batch of configurations, called a *q-batch*, is run in parallel. The resulting AF is referred to as *qCEHVI*. Despite its sample efficiency, the PoF term in qCEHVI is estimated using a combination of Monte Carlo (MC) over fantasy samples (Wilson et al., 2017; Balandat et al., 2020) and soft operators for differentiability (Daulton et al., 2020), thereby introducing bias and a tunable parameter.

Although naïve marginalization can, in principle, eliminate both the sigmoid relaxation and its temperature parameter from qCEHVI, it fails to converge to the exact AF value for batch size $Q > 1$. In this paper, we first analyze that the potential bias of the AF value by naïve marginalization attributes to the candidate strongly correlated with spatially concentrated running points. Our analysis

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

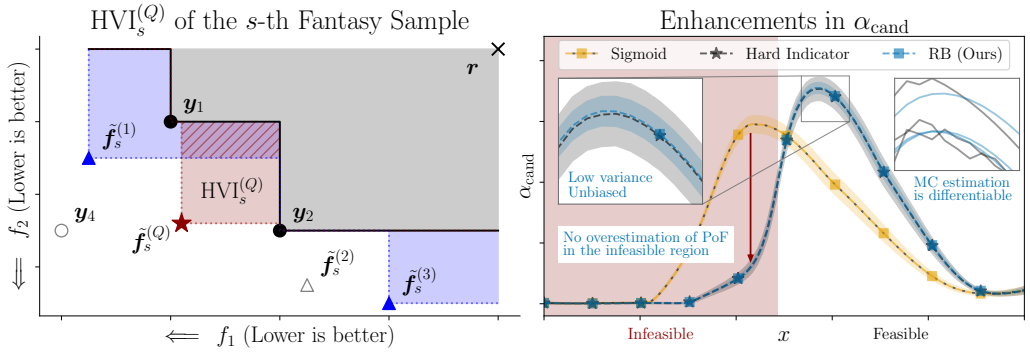


Figure 1: **Left**: Visualization of the HVIs HVI$_s^{(Q)}$ and HVI($\boldsymbol{f}_{\Omega,s}^{(1:Q-1)}|Y_{N,\Omega}$), computed from the candidate fantasy sample (★) and the s-th feasible fantasy sample (▲) respectively, against the observed Pareto solutions (●). The gray hollow objects denote infeasibility. Only the red shaded area changes by $\boldsymbol{x}^{(Q)}$. HVI$_s^{(Q)}$ obtains the hatched area when $\tilde{\boldsymbol{f}}_s^{(1)}$ is infeasible—the error source of the naïve approach discussed in Section 3.1. **Right**: Conceptual AF comparison between our approach and the sigmoid approach against the ground truth (black dashed). The light-shaded bands represent the standard deviation of the MC estimations of α_{cand}. The **left inset** shows that our approach matches the hard indicator mean, which is the asymptotic ground truth. Meanwhile, the sigmoid approach overestimates in the infeasible region, which can hurt performance (Section 3.3). The **right inset** shows that our MC estimation is differentiable unlike ground truth.

shows that the failure arises because naïve marginalization discards the dependence between the candidate's HVI and feasibility induced by the running-point fantasy samples. Given the analysis, we propose to apply Rao Blackwellization (RB) that conditions the candidate's PoF on these fantasy samples to preserve the dependence, while removing the relaxed feasibility indicator. We prove that the RB estimator of qCEHVI is unbiased, low-variance, and pathwise differentiable, which facilitates gradient-based AF optimization, as visualized in Figure 1 (**Right**).

Beyond these theoretical guarantees, our formulation yields empirical advantages. We first confirm on a synthetic problem that our approach resolves the overestimation issue for a large batch size. Furthermore, we find that the sigmoid relaxation bias hinders efficient search near the feasibility boundary. In contrast, our unbiased estimator enables aggressive search near the feasibility boundary stably, a critical factor for overall optimization performance (Picheny, 2014; Liu et al., 2022; Tian et al., 2024). Notably, our formulation, combined with *sequential greedy* batch selection and independence among objectives and constraints—both widely used in practice (Gardner et al., 2014; Hernández-Lobato et al., 2016; Eriksson & Poloczek, 2021; Watanabe & Hutter, 2023; Wilson et al., 2018; Daulton et al., 2020)—is *computable without numerical integration*. Finally, the real-world design experiments demonstrate that our method delivers stable performance compared to the qCEHVI with the sigmoid relaxation, whose efficacy depends on the temperature parameter.

In summary, our contributions are the following: (1) We show that the naïve approach, exact at $Q = 1$, can become biased for $Q > 1$ due to the dependence induced by running-point fantasy samples, and characterize its failure mode. (2) We derive an unbiased, low-variance, pathwise-differentiable estimator of qCEHVI that preserves the dependence via conditional PoF. (3) We show that removing the sigmoid relaxation improves sample efficiency near the feasibility boundary. (4) We demonstrate the RB approach's stability across real-world benchmark problems.

2 Background & Related Work

2.1 Problem Setup & Notations

We consider M black-box objectives to be minimized, $f(\boldsymbol{x}) := (f_1(\boldsymbol{x}), \dots, f_M(\boldsymbol{x}))$, and L black-box constraints, $g(\boldsymbol{x}) := (g_1(\boldsymbol{x}), \dots, g_L(\boldsymbol{x})) \preceq \boldsymbol{0}_L$, over a search space $\mathcal{X} \subseteq \mathbb{R}^d$, where $\preceq$ denotes the elementwise order. Constraint thresholds are absorbed into $\boldsymbol{g}$, so the feasible region is $\Omega := \{\boldsymbol{x} \in \mathcal{X} \mid g(\boldsymbol{x}) \preceq \boldsymbol{0}_L\}$. After N evaluations, the history is $\mathcal{D}_N := \{(\boldsymbol{x}_n, \boldsymbol{y}_n, \boldsymbol{c}_n)\}_{n=1}^N$, where $\boldsymbol{y}_n = f(\boldsymbol{x}_n) + \boldsymbol{\varepsilon}_{f,n} \in \mathbb{R}^M$ and $\boldsymbol{c}_n = g(\boldsymbol{x}_n) + \boldsymbol{\varepsilon}_{g,n} \in \mathbb{R}^L$ are possibly noisy objective and constraint observations. Note, however, that a *large* noise setup, i.e., qCNEHVI (Daulton et al., 2021), is

beyond the scope of our study. Denote the feasible objective set as $Y_{N,\Omega} \coloneqq \{\boldsymbol{y}_n \in Y_N \mid \boldsymbol{c}_n \preceq \mathbf{0}_L\}$ where $Y_N \coloneqq \{\boldsymbol{y}_n\}_{n=1}^N$. Given a finite reference point $\boldsymbol{r} \in \mathbb{R}^M$ that is componentwise worse than the objective vectors of interest, the dominated feasible hypervolume (HV) of a finite set $Y_{N,\Omega}$ is $V(Y_{N,\Omega}, \boldsymbol{r}) \coloneqq \mathrm{vol}(\bigcup_{\boldsymbol{y} \in Y_{N,\Omega}} [\boldsymbol{y}, \boldsymbol{r}))$. We model the latent vectors $(\boldsymbol{f}, \boldsymbol{g})$ with a Gaussian process (GP) posterior conditioned on $\mathcal{D}_N$. Unless stated otherwise, all probabilities and expectations below are taken with respect to this posterior, and the conditioning on $\mathcal{D}_N$ is suppressed. We further denote a batch by $\boldsymbol{x}^{(1:Q)} \coloneqq (\boldsymbol{x}^{(1)}, \ldots, \boldsymbol{x}^{(Q)}) \in \mathcal{X}^Q$ for q-batch BO and its latent vectors by $\boldsymbol{f}^{(q)} \coloneqq f(\boldsymbol{x}^{(q)})$ and $\boldsymbol{g}^{(q)} \coloneqq g(\boldsymbol{x}^{(q)})$ for $q \in [Q]$. The q-batch setup selects a batch of Q points per iteration. Evaluation of the selected points begins asynchronously, without waiting for the entire batch to be assembled. While under evaluation, these are referred to as *running* points.

2.2 Batched Constrained Multi-Objective Bayesian Optimization

Given the history $\mathcal{D}_N$, BO iteratively selects the next configuration $\boldsymbol{x}_{\mathrm{opt}} \in \mathrm{argmax}_{\boldsymbol{x} \in \mathcal{X}} \alpha(\boldsymbol{x})$ where α is the AF, most commonly drawn from the expected improvement family (Ament et al., 2023). The multi-objective variant with $Q = 1$ is the EHVI, $\mathbb{E}[\mathrm{HVI}(\{f(\boldsymbol{x})\}|Y_N)]$ where $\mathrm{HVI}(A'|A) \coloneqq V(A \cup A', \boldsymbol{r}) - V(A, \boldsymbol{r})$ is the HV improvement (HVI). Under outcome constraints, the HVI is multiplied by the feasibility indicator as $\mathrm{HVI}(\{f(\boldsymbol{x})\}|Y_{N,\Omega})\mathbb{I}[g(\boldsymbol{x}) \preceq \mathbf{0}_L]$, so that an infeasible point contributes no improvement (Gardner et al., 2014; Gelbart et al., 2014; Feliot et al., 2017; Daulton et al., 2020). CEHVI averages this quantity over the posterior, namely $\alpha(\boldsymbol{x}) \coloneqq \mathbb{E}[\mathrm{HVI}(\{f(\boldsymbol{x})\}|Y_{N,\Omega})\mathbb{I}[g(\boldsymbol{x}) \preceq \mathbf{0}_L]]$, to select the next point for evaluation. Although Deng et al. (2025) proposed an analytical form of qEHVI, introducing either the log-transformation (Ament et al., 2023) or outcome constraints breaks this exact formulation. If objectives and constraints are mutually independent, the AF factorizes as $\alpha(\boldsymbol{x}) = \mathbb{E}[\mathrm{HVI}(\{f(\boldsymbol{x})\}|Y_{N,\Omega})]\mathbb{P}[g(\boldsymbol{x}) \preceq \mathbf{0}_L]$. The independence among constraints is a common choice in prior work (Gardner et al., 2014; Gelbart et al., 2014; Hernández-Lobato et al., 2015; Eriksson & Poloczek, 2021; Watanabe & Hutter, 2023), enabling the joint PoF to further factorize $\mathbb{P}[g(\boldsymbol{x}) \preceq \mathbf{0}_L] = \prod_{l=1}^L \mathbb{P}[g_l(\boldsymbol{x}) \leq 0]$. Zhang et al. (2024) recently recommended independent modeling, as joint modeling did not provide a meaningful performance advantage in their study.

For $Q > 1$, qCEHVI is computed as $\mathbb{E}[\mathrm{HVI}(\boldsymbol{f}_\Omega^{(1:Q)}|Y_{N,\Omega})]$ where $\boldsymbol{f}_\Omega^{(1:q')} \coloneqq \{\boldsymbol{f}^{(q)} \mid \boldsymbol{g}^{(q)} \preceq \mathbf{0}_L, q \in [q']\}$. Daulton et al. (2021) formulated qCEHVI specifically for the sequential greedy scenario where $\boldsymbol{x}^{(1:Q-1)}$ are fixed running points and $\boldsymbol{x}^{(Q)}$ is the candidate to be selected. The full batch HVI then decomposes as

$$\mathrm{HVI}(\boldsymbol{f}_\Omega^{(1:Q)} \mid Y_{N,\Omega}) = \underbrace{\mathrm{HVI}(\boldsymbol{f}_\Omega^{(1:Q-1)} \mid Y_{N,\Omega})}_{=\text{constant w.r.t. } \boldsymbol{x}^{(Q)}} + \underbrace{\mathrm{HVI}(\{\boldsymbol{f}^{(Q)}\} \mid Y_{N,\Omega} \cup \boldsymbol{f}_\Omega^{(1:Q-1)})}_{=\mathrm{HVI}^{(Q)}} \underbrace{\mathbb{I}[\boldsymbol{g}^{(Q)} \preceq \mathbf{0}_L]}_{=\mathbb{I}_\Omega^{(Q)}}, \tag{1}$$

where the first term is the running-only HVI and the second is the candidate's actual contribution. Here, $\mathrm{HVI}^{(Q)}$ is the additional HVI that the candidate would contribute if feasible, and $\mathbb{I}_\Omega^{(Q)}$ is its feasibility indicator. Since the first term is constant with respect to $\boldsymbol{x}^{(Q)}$, only the candidate's EHVI needs to be maximized. Figure 1 (**Left**) intuitively visualizes the formulation. Using this formulation, Daulton et al. (2021) enabled calculating qCEHVI in polynomial time instead of the original exponential time (Daulton et al., 2020). Empirical evidence indicates that sequential greedy achieves better sample efficiency than joint optimization, as solving lower-dimensional steps avoids a high-dimensional problem (Wilson et al., 2018; Wada & Hino, 2019; Daulton et al., 2020).

Empirically, the candidate's EHVI $\alpha_{\mathrm{cand}} \coloneqq \mathbb{E}[\mathrm{HVI}^{(Q)}\mathbb{I}_\Omega^{(Q)}]$ is calculated using the MC estimator $\frac{1}{S}\sum_{s=1}^S \mathrm{HVI}_s^{(Q)}\mathbb{I}_{\Omega,s}^{(Q)}$, where $\mathrm{HVI}_s^{(Q)}$ and $\mathbb{I}_{\Omega,s}^{(Q)}$ are the additional HVI and feasibility indicator computed from the candidate's s-th fantasy sample from the posterior, $\tilde{\boldsymbol{f}}_s^{(Q)}$ and $\tilde{\boldsymbol{g}}_s^{(Q)}$. Since the hard indicator $\mathbb{I}_{\Omega,s}^{(Q)}$ has zero derivatives with respect to the sampled constraint values almost everywhere, Daulton et al. (2021) replaced it with the sigmoid weight $w_{\tau,s}^{(Q)} \coloneqq \prod_{l=1}^L \varsigma(-\tilde{g}_{s,l}^{(Q)}/\tau)$, where $\varsigma(t) \coloneqq (1 + \exp(-t))^{-1}$ is a sigmoid function and $\tau > 0$ controls the sharpness of the approximation. This gives the relaxed MC estimator of the candidate's EHVI $\frac{1}{S}\sum_{s=1}^S \mathrm{HVI}_s^{(Q)} w_{\tau,s}^{(Q)}$. Despite its differentiability, the soft PoF term $w_{\tau,s}^{(Q)}$ is a biased estimator of the hard indicator $\mathbb{I}_{\Omega,s}^{(Q)}$, and the temperature parameter τ needs to be tuned. Following Ament et al. (2023), who discuss the performance advantages of the log-transformed AF, we optimize $\log \alpha_{\mathrm{cand}}$ during the AF optimization, but α_{cand} consistently represents the raw value throughout this paper. By contrast, for $Q > 1$, Daulton et al. (2020); Ament et al. (2023) reported no known analytical expressions and instead used joint MC estimation. Section 3 shows that the candidate's PoF can nevertheless be integrated out exactly.

3 MARGINALIZATION CAN ELIMINATE SIGMOID FROM qCEHVI

In this section, we first discuss the elimination of the sigmoid via marginalization. However, our analysis reveals that naïve marginalization can cause overestimation for large batches with correlated running points. We then describe why this overestimation matters, and propose an unbiased estimation using Rao-Blackwellization (RB). Finally, we show that our method resolves a pathological behavior near the feasibility boundary caused by the sigmoid relaxation bias, and describe how this property helps during optimization. Note that the theoretical discussion uses sequential greedy for notational simplicity, but our results generalize without the restriction as stated in Section 3.4.

3.1 MARGINAL POF FAILS UNDER BATCH-INDUCED DEPENDENCE

Daulton et al. (2020) computed the expected candidate's HVI α_{cand} under $Q = 1$ and independence between constraints and objectives by:

$$\alpha_{\mathrm{cand}}(\boldsymbol{x}) = \mathbb{E}[\mathrm{HVI}(\{\boldsymbol{f}\} \mid Y_{N,\Omega})\mathbb{I}[\boldsymbol{g} \preceq \mathbf{0}_L]] = \mathbb{E}[\mathrm{HVI}(\{\boldsymbol{f}\} \mid Y_{N,\Omega})]\mathbb{P}[\boldsymbol{g} \preceq \mathbf{0}_L], \tag{2}$$

while they proposed the sigmoid approach for $Q > 1$ due to the exact form's analytical intractability. Naïvely, this formula is extended to $\mathbb{E}[\mathrm{HVI}^{(Q)}]\mathbb{E}[\mathbb{I}_\Omega^{(Q)}] = \mathbb{E}[\mathrm{HVI}^{(Q)}]p_{\mathrm{marg}}^{(Q)}$ where $p_{\mathrm{marg}}^{(Q)} \coloneqq \mathbb{P}[\boldsymbol{g}^{(Q)} \preceq \mathbf{0}_L]$ is the marginal PoF. However, we show that this intuition does not generalize to $Q > 1$.

Proposition 1 (Marginal-PoF Bias)**.** *For any fixed batch,* $\hat{\alpha}_{\mathrm{marg},S} \coloneqq p_{\mathrm{marg}}^{(Q)} \frac{1}{S}\sum_{s=1}^{S} \mathrm{HVI}_s^{(Q)}$ *has the following bias:*

$$\mathbb{E}[\hat{\alpha}_{\mathrm{marg},S}] - \alpha_{\mathrm{cand}} = -\mathrm{Cov}\big[\mathrm{HVI}^{(Q)},\ \mathbb{I}_\Omega^{(Q)}\big]. \tag{3}$$

Proof Sketch. As $p_{\mathrm{marg}}^{(Q)} = \mathbb{E}[\mathbb{I}_\Omega^{(Q)}]$ and $\mathbb{E}[\hat{\alpha}_{\mathrm{marg},S}] = \mathbb{E}[\mathrm{HVI}^{(Q)}]p_{\mathrm{marg}}^{(Q)}$, subtracting $\alpha_{\mathrm{cand}} = \mathbb{E}[\mathrm{HVI}^{(Q)}\mathbb{I}_\Omega^{(Q)}]$ gives the identity. □

See Theorem 12 and Remark 15 for a more formal treatment. Proposition 1 shows that the covariance between $\mathrm{HVI}^{(Q)}$ and $\mathbb{I}_\Omega^{(Q)}$ induces bias in the estimation. As stated in Proposition 19 and Remark 20, *overestimation* is typical when the candidate lies near the running points. More concretely, when a running point is feasible, i.e., $\boldsymbol{g}^{(q)} \preceq \mathbf{0}_L$, and $\boldsymbol{f}^{(Q)} \simeq \boldsymbol{f}^{(q)}$, $\mathrm{HVI}^{(Q)}$ is significantly reduced. For example, in Figure 1 (**Left**), since $\boldsymbol{f}_s^{(1)}$ is feasible, the hatched area does not contribute to $\mathrm{HVI}^{(Q)}$. The contribution reduction is more significant when $\boldsymbol{f}_s^{(Q)} \simeq \boldsymbol{f}_s^{(1)}$. This example illustrates that the optimal candidate depends on the constraint latent vectors at the running points. Proposition 14 details that $Q = 1$ is a sufficient condition for zero marginal PoF bias. We discuss the solution to this issue and demonstrate the fix on a synthetic problem.

3.2 RAO-BLACKWELLIZATION—ELIMINATING SIGMOID WITHOUT BIAS

Proposition 1 revealed that the bias arises from the dependence between the candidate's HVI and PoF induced by the running-point fantasy samples. Denote $Z \coloneqq (\boldsymbol{f}^{(1:Q)}, \boldsymbol{g}^{(1:Q-1)})$ as the latent vectors except the candidate's constraint latent vector $\boldsymbol{g}^{(Q)}$, and assume finiteness of an arbitrary moment $\mathbb{E}[(\mathrm{HVI}^{(Q)})^k] < \infty$. We apply *Rao-Blackwellization* (RB) to obtain the *conditional* PoF $p_{\mathrm{RB}}^{(Q)} \coloneqq \mathbb{P}[\boldsymbol{g}^{(Q)} \preceq \mathbf{0}_L | Z]$ and take the average of $p_{\mathrm{RB},s}^{(Q)}\mathrm{HVI}_s^{(Q)}$ for each joint fantasy sample Z_s:

$$\hat{\alpha}_{\mathrm{RB},S} \coloneqq \frac{1}{S}\sum_{s=1}^{S} \mathrm{HVI}_s^{(Q)} p_{\mathrm{RB},s}^{(Q)}, \tag{4}$$

where $p_{\mathrm{RB}}^{(Q)}$ is also Gaussian, thanks to the closed-form normality of Gaussian distributions. This formulation integrates out the candidate's constraint latent vector, while the running points' fantasy sample remains to determine $\mathrm{HVI}_s^{(Q)}$. Each HVI is therefore weighted by the candidate's conditional PoF, preserving the dependence lost when using the marginal PoF. The following theorem shows that Eq. (4) yields an unbiased sample average, unlike $\hat{\alpha}_{\mathrm{marg},S}$, while retaining the differentiability that the hard indicator approach lacks and achieving lower variance than it:

Theorem 2 (RB Is a Differentiable Unbiased Estimator of α_{cand} with Lower Variance)**.** *For any fixed batch, the following holds:*

(i) ***Exactness****: the MC estimator in Eq. (4) satisfies* $\mathbb{E}[\hat{\alpha}_{\mathrm{RB},S}] = \alpha_{\mathrm{cand}}$,

(ii) ***Differentiability****:* $\hat{\alpha}_{\mathrm{RB},S}$ *is locally Lipschitz and differentiable almost everywhere, with*

$$\nabla_{\boldsymbol{x}^{(Q)}}\hat{\alpha}_{\mathrm{RB},S} = \frac{1}{S}\sum_{s=1}^{S}\left(p_{\mathrm{RB},s}^{(Q)}\nabla_{\boldsymbol{x}^{(Q)}}\mathrm{HVI}_s^{(Q)} + \mathrm{HVI}_s^{(Q)}\nabla_{\boldsymbol{x}^{(Q)}}p_{\mathrm{RB},s}^{(Q)}\right), \tag{5}$$

(iii) ***Variance reduction****:* $\mathbb{V}[\mathrm{HVI}^{(Q)}\mathbb{I}_\Omega^{(Q)}] - \mathbb{V}[\mathrm{HVI}^{(Q)}p_{\mathrm{RB}}^{(Q)}] = \mathbb{E}[(\mathrm{HVI}^{(Q)})^2 p_{\mathrm{RB}}^{(Q)}(1 - p_{\mathrm{RB}}^{(Q)})] \geq 0$.

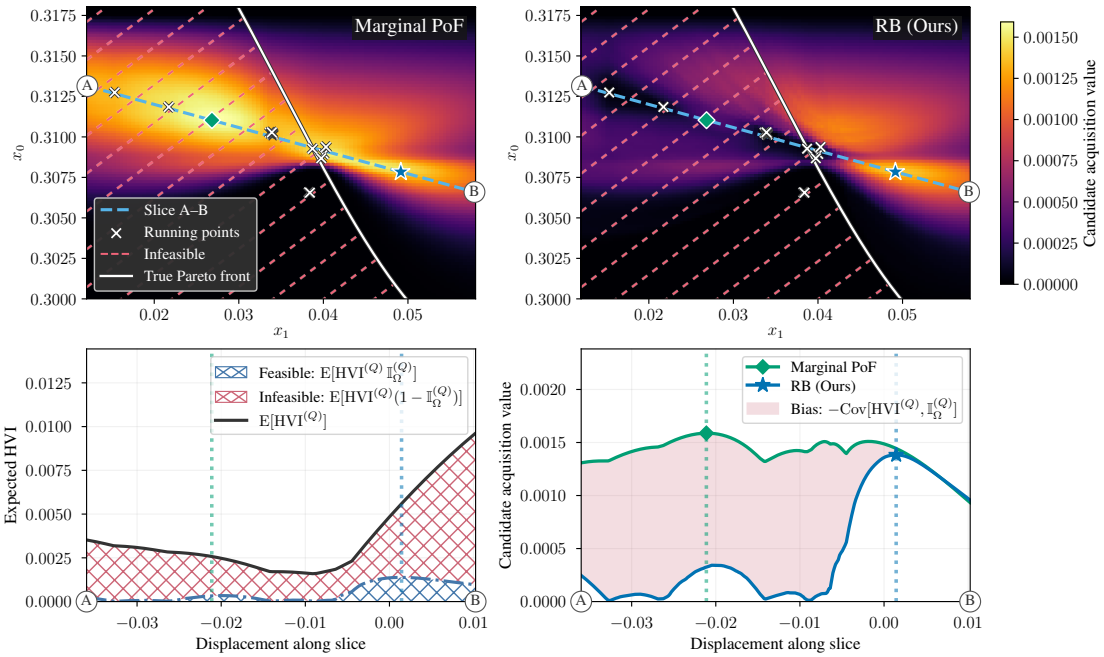


Figure 2: Marginal-PoF bias on TNK (Tanaka et al., 1995) with $Q = 16$ and 15 running points (shown by white $\times$). The slice A-B connects two maxima marked by ★ for $\hat{\alpha}_{\mathrm{RB},S}$ and ◆ for $\hat{\alpha}_{\mathrm{marg},S}$. **Top:** AF landscapes (brighter indicates a higher value). $\hat{\alpha}_{\mathrm{marg},S}$ (**Top Left**) exhibits high AF values near the region where the running points are concentrated, while $\hat{\alpha}_{\mathrm{RB},S}$ (**Top Right**) does not. **Bottom Left:** Infeasible and feasible EHVIs along the slice. The vertical dashed lines represent the maxima locations (★, ◆) in the top panels. The overestimation is caused by incorrectly counting part of the infeasible EHVI as feasible EHVI. **Bottom Right:** AF values along the slice. The blue line is identical to the feasible EHVI line in **Bottom Left**. Roughly speaking, $\hat{\alpha}_{\mathrm{marg},S}$ incorrectly weights the infeasible EHVI near ◆ by overestimating the PoF there. Given infeasible fantasy samples at the running points, $\hat{\alpha}_{\mathrm{RB},S}$ calibrates the PoF downward, suppressing the AF value near ◆. The shaded area represents the bias $-\mathrm{Cov}[\mathrm{HVI}^{(Q)}, \mathbb{I}_\Omega^{(Q)}]$.

Proof Sketch. (i) The additional HVI is a function of Z, and $\mathbb{E}[\mathbb{I}_\Omega^{(Q)} \mid Z] = p_{\mathrm{RB}}^{(Q)}$. Hence,

$$\mathbb{E}[\mathrm{HVI}^{(Q)}\mathbb{I}_\Omega^{(Q)} \mid Z] = \mathrm{HVI}^{(Q)}\mathbb{E}[\mathbb{I}_\Omega^{(Q)} \mid Z] = \mathrm{HVI}^{(Q)} p_{\mathrm{RB}}^{(Q)}. \tag{6}$$

Taking expectations gives $\mathbb{E}[\mathrm{HVI}^{(Q)} p_{\mathrm{RB}}^{(Q)}] = \alpha_{\mathrm{cand}}$, and linearity of expectation yields the result.

(ii) The conditional PoFs and Cholesky-generated fantasy samples are continuously differentiable. Since fixed fantasy fronts make each HVI locally Lipschitz in $\boldsymbol{x}^{(Q)}$, the result follows from Rademacher's theorem and the product rule.

(iii) Conditional on Z, the additional HVI is fixed and $\mathbb{I}_\Omega^{(Q)}$ is Bernoulli with success probability $p_{\mathrm{RB}}^{(Q)}$. Hence, we have $\mathbb{V}\big[\mathrm{HVI}^{(Q)}\mathbb{I}_\Omega^{(Q)} \mid Z\big] = (\mathrm{HVI}^{(Q)})^2 p_{\mathrm{RB}}^{(Q)}(1 - p_{\mathrm{RB}}^{(Q)})$. The law of total variance, together with the conditional-expectation identity in Eq. (6), gives the result. □

More formal treatments for Results (i)-(iii) are discussed in Theorems 9, 10 and 17, respectively. Theorem 2 showed that qCEHVI can be correctly estimated even without a sigmoid for $Q > 1$, while keeping the differentiability. Moreover, our approach reduces variance compared to the hard indicator-based approach as a byproduct.

We now empirically walk through when the naïve approach is pathological, and how RB alleviates the issue. Figure 2 shows the scenario where the marginal-PoF bias degrades optimization performance for $Q = 16$. The top panels visualize the AF landscapes for the naïve approach (**Left**) and our approach (**Right**) with the shared set of running points represented by the white cross marks. This example reproduces spatially concentrated running points, commonly observed for large batch sizes. The pitfall here is that the naïve approach incorrectly yields a high AF value in the infeasible region due to the overestimation. Figure 2 (**Bottom Left**) shows EHVI (solid black) and its breakdown. Since we need to consider the constraints along with EHVI, our interest is the blue shaded area. Although our approach correctly computes its unbiased estimate, the naïve approach ends up with the overestimation shaded red in Figure 2 (**Bottom Right**).

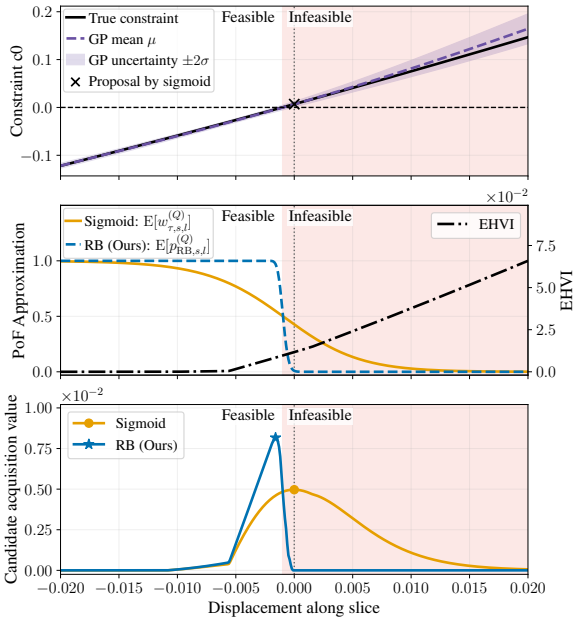


Figure 3: AF comparison with $Q = 16$ between the RB and sigmoid approaches on TNK (Tanaka et al., 1995). The slice is defined by connecting the maxima for each approach (★: RB, ●: sigmoid). **Top**: The true constraint (black solid) and its estimation by a GP (purple shade: $\pm 2\sigma$). The feasibility boundary starts when the true constraint crosses the constraint threshold line (black dashed). **Middle**: Expected PoF approximations by a sigmoid $\mathbb{E}[w^{(Q)}_{\tau,s,l}]$ and RB $\mathbb{E}[p^{(Q)}_{\mathrm{RB},s,l}]$, with EHVI shown on the right axis. The approximation by a sigmoid remains positive well beyond the feasibility boundary, whereas that by RB is nearly zero. EHVI (black solid) increases towards the infeasible region. **Bottom**: AF values along the slice (sigmoid, RB). The sigmoid approach peaks in the infeasible region, driven by the higher EHVI there. By contrast, the RB approach peaks in the near-boundary feasible region with nearly zero AF in the infeasible region.

The bias essentially comes from disregarding a fantasy sample's feasibility. For example, the shaded area in Figure 1 (**Left**) cuts off a part of the potential candidate's EHVI, if $\tilde{\boldsymbol{f}}^{(1)}_s$ is feasible. Since $\tilde{\boldsymbol{f}}^{(1)}_s$ and $\tilde{\boldsymbol{f}}^{(Q)}_s$ are close to each other, the candidate point is likely feasible. Although the potential candidate's EHVI becomes larger when running points nearby are infeasible, the candidate is likely infeasible at such a point, leading to $\mathbb{I}^{(Q)}_\Omega = 0$. Disregarding this intuition is indeed the source of the overestimation. Mathematically speaking, a fantasy sample is drawn from a joint GP posterior, inducing high correlation in the feasibility among concentrated points. Therefore, the posterior at the candidate point has to consider the correlation with the running points. The RB approach does so by conditioning the fantasy sample, identifying the maximum feasible HVI near ★. By contrast, since $\hat{\alpha}_{\mathrm{marg},S}$ scales $\mathrm{HVI}^{(Q)}$ by $p^{(Q)}_{\mathrm{marg}}$, which does not change regardless of which running point is infeasible, it cannot tell the two maxima apart and inherits the infeasible EHVI's magnitude near ◆, changing its maximizer as explained in Remark 13. The shaded bias term $-\mathrm{Cov}[\mathrm{HVI}^{(Q)}, \mathbb{I}^{(Q)}_\Omega]$ in Figure 2 (**Bottom Right**) confirms this directly. Remark 15 discusses that the upper bound of the marginal-PoF bias is large at the region where the marginal PoF or $\mathrm{HVI}^{(Q)}$ is uncertain. Our approach provides a robust estimator that stabilizes optimization under uncertainty.

3.3 Is the RB Approach Well-Behaved?—Accurate Infeasibility Awareness

As discussed earlier, our method uses an unbiased PoF estimator, while the sigmoid approach does not. This implies that the sigmoid approach can overestimate the promise of candidates, even if the GP confidently predicts infeasibility. Let $\mu_{l,s}$ and $\sigma_{l,s}$ denote the conditional mean and standard deviation of $g^{(Q)}_l$ given Z_s.

Proposition 3 (Sigmoid Does Not Vanish for Highly Confident Infeasibility)**.** *Fix $\tau > 0$ and the means $\mu_{l,s}$, with $\mu_{k,s} > 0$ for some k. As $\sigma_{l,s} \to 0^+$ for all l, we have*

$$p^{(Q)}_{\mathrm{RB},s} = \mathbb{P}[\boldsymbol{g}^{(Q)} \preceq \mathbf{0}_L \mid Z_s] \to 0,\ w^{(Q)}_{\tau,s} \xrightarrow{P} \prod_{l=1}^{L} \varsigma\left(-\frac{\mu_{l,s}}{\tau}\right) > 0, \tag{7}$$

where $\xrightarrow{P}$ denotes convergence in probability.

Proof Sketch. We have $0 \le p^{(Q)}_{\mathrm{RB},s} \le \mathbb{P}[g^{(Q)}_k \le 0 \mid Z_s] = \Phi(-\mu_{k,s}/\sigma_{k,s}) \to 0$, where Φ is the standard normal cumulative distribution function. As $\sigma_{l,s} \to 0^+$ for all l, the samples $\tilde{g}^{(Q)}_{s,l}$ converge jointly in probability to $\mu_{l,s}$, so continuity and positivity of the sigmoid give the second limit. □

Thus, when $\mathrm{HVI}^{(Q)}_s > 0$ for some s, $\hat{\alpha}_{\mathrm{RB},S} = \frac{1}{S}\sum_{s=1}^{S} \mathrm{HVI}^{(Q)}_s p^{(Q)}_{\mathrm{RB},s} \to 0$, whereas the sigmoid estimator $\hat{\alpha}_{\mathrm{sig},S} := \frac{1}{S}\sum_{s=1}^{S} \mathrm{HVI}^{(Q)}_s w^{(Q)}_{\tau,s}$ remains positive. Informally, a sufficiently large HVI can remain attractive under the sigmoid approach even when the GP confidently predicts a constraint violation, whereas our approach discards such points.

Figure 3 intuitively illustrates this behavior. In this problem setup, EHVI grows towards the infeasible region, necessitating accurate feasibility awareness. Indeed, this problem setup exemplifies a widely recognized problem class where constrained optima lie near feasibility boundaries (Picheny,

2014; Liu et al., 2022; Tian et al., 2024). In general, observing constraint violations reduces $\sigma_{l,s}$ near the evaluated points while establishing $\mu_{l,s} > 0$. Figure 3 (**Top**) visualizes this phenomenon, enabling $p_{\mathrm{RB}}^{(Q)}$ to approximate the boundary accurately. The GP correctly predicts $\mu_{l,s} > 0$ in the infeasible region with the thin standard deviation band, i.e., $\sigma_{l,s} \simeq 0^+$. Equation (7) enforces $\hat{\alpha}_{\mathrm{RB},S} \to 0$ in the infeasible region under this scenario. Therefore, the RB approach sharply drops at the feasibility boundary in Figure 3 (**Bottom**). Interestingly, this behavior can be explained by strong repulsive gradients against highly confident infeasibility as discussed in Remark 24. In contrast, the AF of the sigmoid approach remains positive and incorrectly attains its optimum in the infeasible region, meaning the GP's infeasibility confidence is not fully reflected in the AF. Infeasible candidates with large HVI become attractive with this positive residual, leading to repeated infeasible evaluations even after the GP has fully accounted for the feasibility boundaries. Although $\tau \to 0^+$ mitigates the issue, the sigmoid relaxation approaches the hard indicator $\mathbb{I}_{\Omega}^{(Q)}$, so the increasingly sharp transition makes the AF optimization harder while still requiring parameter tuning.

3.4 Practical Considerations

Our implementation follows sequential greedy and independent GPs due to the practical advantages discussed in Section 2. Eq. (4) then becomes the following closed form:

$$\hat{\alpha}_{\mathrm{RB},S} = \frac{1}{S} \sum_{s=1}^{S} \mathrm{HVI}_s^{(Q)} \prod_{l=1}^{L} \Phi(-\mu_{l,s}/\sigma_{l,s}), \tag{8}$$

where $\mu_{l,s}$ and $\sigma_{l,s}$ are the conditional moments in Section 3.3. The sigmoid approach instead computes $\varsigma(-\tilde{g}_{s,l}^{(Q)}/\tau)$ using a fantasy sample from the same moments. Since the two estimators share the GP posterior and the Cholesky factors, both require $O(SL)$, and the inference cost difference comes solely from evaluating Φ versus ς. We therefore do not report a separate runtime comparison in Section 4. Recall that Theorem 2 uses neither assumption, so exactness, variance reduction, and differentiability are preserved even when either is dropped. Dropping independence turns $p_{\mathrm{RB}}^{(Q)}$ into an L-dimensional Gaussian CDF, and dropping sequential greedy makes the exact AF enumerate $2^Q - 1$ batch subsets (Daulton et al., 2020) where RB weights each term by the conditional probability that all its points are feasible (Remark 25).

4 Experiments

Finally, we benchmark our method against widely used baselines on both synthetic and real-world design configuration problems. Due to the equivalence among RB, marginal PoF, and sigmoid at $Q = 1$, all three yield perfectly identical results with a fixed seed. Therefore, the goal of this experiment is to not demonstrate the performance superiority of the RB approach among existing state-of-the-art methods but rather to justify replacing the sigmoid approach with the tuning-free RB approach. To this end, our evaluation focuses on whether the RB approach is stable and whether its performance is indistinguishable from or superior to the marginal PoF and sigmoid approaches.

4.1 Experimental Setup

Baseline Methods. Since Optuna (Akiba et al., 2019) is widely used in practice for black-box optimization, we chose official methods that support CMO in this package: (1) random search (Bergstra & Bengio, 2012), (2) NSGA-II (Deb et al., 2002), (3) NSGA-III (Deb & Jain, 2014; Jain & Deb, 2013), and (4) TPE (Ozaki et al., 2020; Watanabe & Irie, 2026). In addition, we also compare our method against α_{cand} with the sigmoid relaxation and against α_{cand} with the marginal PoF discussed in Section 3.1. Notice that GP-based BO implementations follow Optuna, which uses qLogCEHVI with $\tau{=}10^{-2}$. This means that our pipeline utilizes sequential greedy and independent GPs, which are recommended in practice, as discussed in Section 2. See Appendix B for more details.

Benchmark Problems. For synthetic functions, we selected TNK (Tanaka et al., 1995) and SRN (Srinivas & Deb, 1994; Garrido-Merchán & Hernández-Lobato, 2020). TNK is a non-linear bi-objective ($M{=}2$) benchmark instance with two non-linear inequality constraints ($L{=}2$) featuring multiple disconnected Pareto fronts over a two-dimensional continuous search space ($d{=}2$) with the feasibility ratio $p_{\mathrm{feas}}{\simeq}5\%$, which is estimated by Monte-Carlo sampling with the sample size of 10^6. SRN is characterized by its strict feasibility ratio ($M{=}2, L{=}2, d{=}2, p_{\mathrm{feas}}{\simeq}0.2\%$). Efficient search on these problems requires accurate infeasibility awareness.

For real-world problems, we used Process Synthesis ($M{=}2, L{=}9, d{=}7, p_{\mathrm{feas}}{<}10^{-4}\%$) (Kocis & Grossmann, 1989), Disc Brake ($M{=}2, L{=}4, d{=}4, p_{\mathrm{feas}}{\simeq}64\%$) (Ray & Liew, 2002), and Multi product Batch Plant ($M{=}3, L{=}10, d{=}10, p_{\mathrm{feas}}{\simeq}0.6\%$) (Dhiman & Kumar, 2018). Process Synthesis

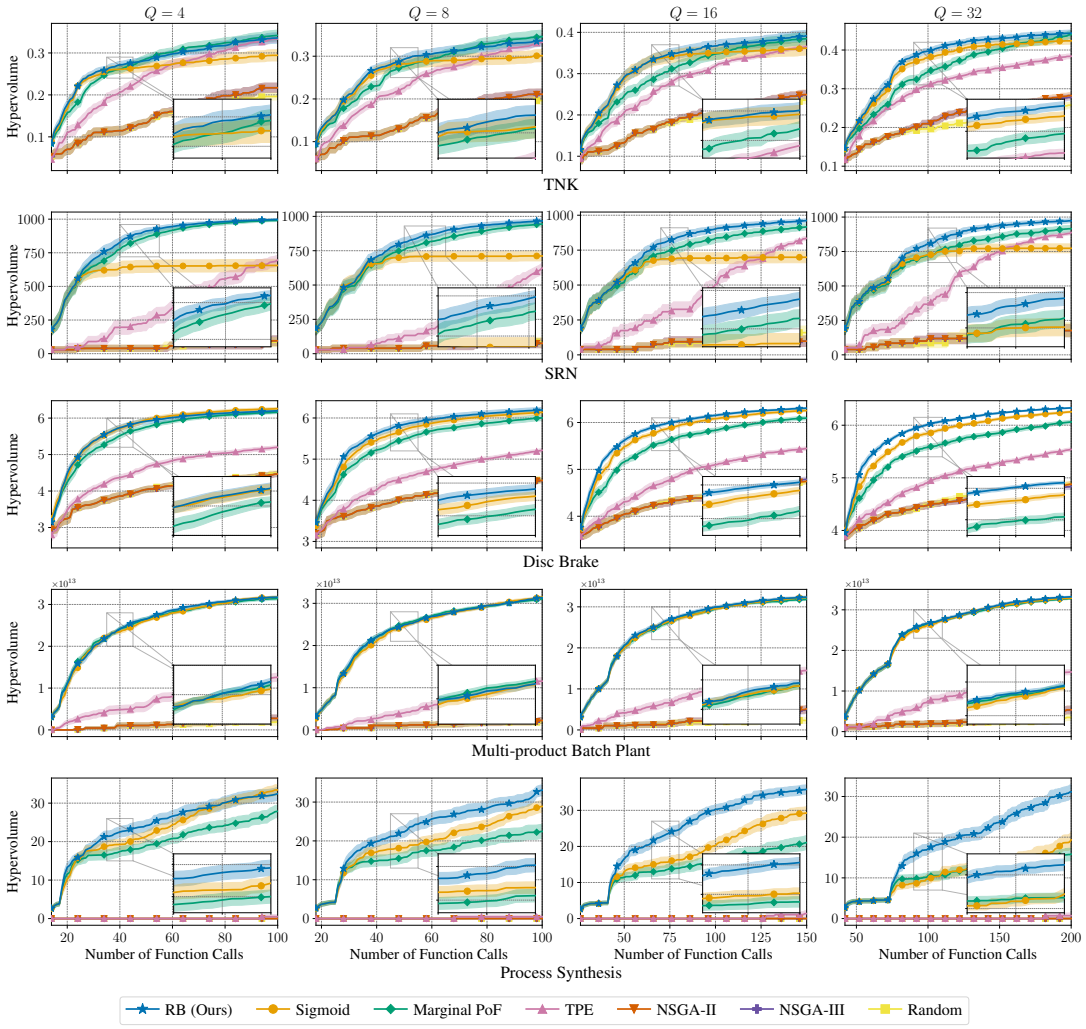


Figure 4: Hypervolume of the feasible observations found versus the number of completed function evaluations (higher is better). Rows show TNK, SRN, Disc Brake, Multi-product Batch Plant, and Process Synthesis; columns correspond to the batch size of $Q \in \{4, 8, 16, 32\}$. The lines and the light-shaded bands denote the mean and the standard error, respectively, both computed over 30 independent runs with different random seeds. Insets magnify early-stage differences. RB (★) achieves high hypervolume stably even at large Q.

aims to optimize the structural layout and operating conditions of a chemical process to minimize deviations from ideal performance targets over a mixed search space with three continuous variables for physical operations and four binary variables for flowsheet configuration. Disc Brake seeks to optimize an automotive disc brake by minimizing its weight and braking time. The design relies on four continuous variables like disc dimensions and engaging force, constrained by thickness, pressure, and temperature. Multi-product Batch Plant focuses on designing a chemical plant that produces multiple products while minimizing costs and overall production time. To simulate asynchronous optimization, we followed the scheme proposed by Watanabe et al. (2024) for correctness.

4.2 Results and Discussion

Figure 4 plots the feasible HV history for these problems. For TNK, Process Synthesis, and Disc Brake, performance of marginal PoF and RB diverges as Q grows. In fact, some marginal PoF runs failed for large Q due to numerical instability of the Cholesky decomposition induced by concentrated suggestions. The results on SRN and Process Synthesis clearly demonstrate the failure of the sigmoid approaches. Despite these results, the three methods including ours share an identical pipeline at $Q = 1$, leading to indistinguishable performance as observed in Multi-product Batch Plant. Since we aim for a tuning-free yet stable method, these results are expected.

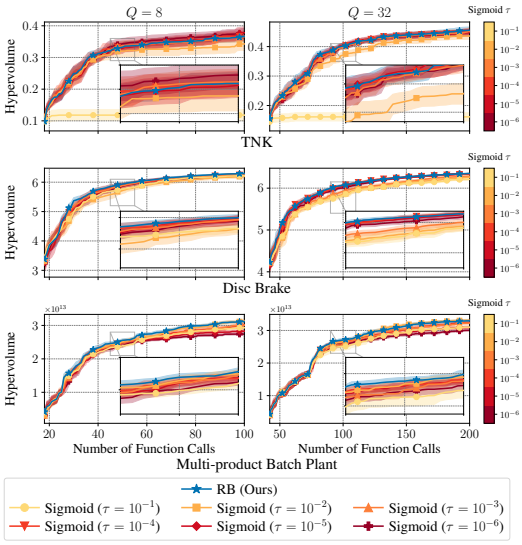


Figure 6: Sensitivity of the sigmoid approach to the temperature $\tau \in \{10^{-1}, 10^{-2}, \ldots, 10^{-6}\}$, compared with RB, which requires no τ. Rows show TNK, Disc Brake, and Multi-product Batch Plant; columns correspond to $Q \in \{8, 32\}$. The lines and the light-shaded bands denote the mean and the standard error, respectively, both computed over 10 independent runs with different random seeds. Multi-product Batch Plant (**Bottom**) yielded worse performance with smaller τ, especially in the late stage, whereas TNK (**Top**) and Disc Brake (**Middle**) performed worse with larger τ.

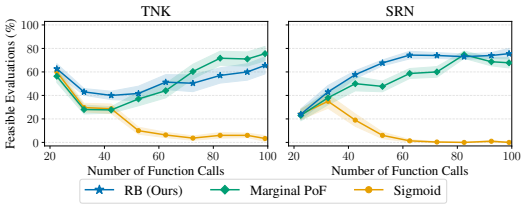


Figure 5: Feasibility ratio on TNK and SRN with Q=8. Ratios are computed over a rolling window of the 10 most recent evaluations. The lines and the light-shaded bands denote the mean and the standard error, respectively, both computed over 30 independent runs with different random seeds. The sigmoid approach's ratio declines, validating Section 3.3's discussion.

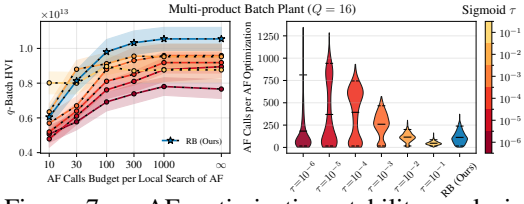


Figure 7: AF optimization stability analysis of small τ on Multi-product Batch Plant with Q=16. Each experiment was performed over 30 independent runs with different random seeds. **Left**: q-batch HVI against completed observations for varied AF optimization budget. Large τ is more robust to a smaller budget, while small τ requires more budget. **Right**: AF calls needed for convergence per local search. Small τ needs more AF calls for convergence.

Figure 5 further illustrates prolonged HV stagnation of the sigmoid approach after the early stage on TNK and SRN. The feasibility ratio of the sigmoid approach gradually declines toward zero thereafter, whereas RB and marginal PoF recover. Repeated violations make GP increasingly confident about infeasibility predictions. Nevertheless, the sigmoid relaxation retains a positive PoF, so candidates with a large HVI can remain attractive despite these predictions. This is consistent with Proposition 3 and Section 3.3.

Figure 6 shows the sensitivity of the sigmoid approach to the temperature parameter τ. The result on TNK shows that a large τ can significantly underperform, making smaller τ promising. By contrast, the result on Multi-product Batch Plant provides the scenario where small $\tau \in \{10^{-3}, \ldots, 10^{-6}\}$ performs significantly worse than RB ($p < 0.05$ by the Wilcoxon signed-rank test). Theoretically speaking, $\mathbb{E}[\hat{\alpha}_{\mathrm{sig},S}] \to \mathbb{E}[\hat{\alpha}_{\mathrm{RB},S}]$ ($\tau \to 0^+$) from Theorem 2, so the performance should be similar for small τ. To further analyze the behavior in this problem, Figure 7 visualizes AF call counts and the obtained q-batch HVI, capped by a maximum AF call budget during the local search. Note that an infinite budget was consistently used in the other experiments to observe the true performance of small τ. For fairness, we collected $\mathcal{D}_N$ (N=200) using TPE for all setups. Figure 7 (**Right**) shows that smaller τ needs more iterations to converge, resulting in poorer performance for small τ in Figure 7 (**Left**) when setting a smaller budget. Note that since gradient-based optimization with τ=10^{-6} did not work well with a poor starting point, some runs ended up stuck at that point, explaining the peak at the small AF calls and highlighting the importance of stable gradients.

5 Conclusion

This paper discussed the pitfalls of eliminating the sigmoid function from qCEHVI. Our empirical and theoretical analyses explained the downsides of the sigmoid and naïve marginalization, and showed our approach resolves these issues. Experiments demonstrated our method performs stably and confirmed our empirical findings generalize to real-world problems, justifying the RB approach.

AI USE STATEMENT

In this work, we used generative AI tools to assist in the writing of proofs, formulate mathematical claims, and assist with translation. We have not used generative AI tools to help develop theoretical models or conceptual frameworks, provide critical ingredients for proving mathematical claims, propose or refine hypotheses, design or provide feedback on research methodology or experiments, implement methods, clean and reformat dataset, support qualitative and thematic data analysis, and interpret results. Also, generating synthetic data sets are not applicable to this work. Additionally, we used generative AI tools to improve readability, create or edit software code (documentation and minor refactoring). We have reviewed all AI-assisted work. All AI-assisted text, code, and documentation, including English-language editing and translation, were thoroughly reviewed by the authors. The authors verified their correctness and appropriateness and substantially revised them as needed before incorporating them into the final work. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

ETHICS STATEMENT

This section is not applicable to this work.

REPRODUCIBILITY STATEMENT

We list the references for each source:

1. Implementation: See "Supplementary Material". The implementation will be publicly available after acceptance.
2. Theoretical Assumptions: Appendix A.1 for all results and Appendix A.6 for the results specifically under the sequential greedy and independent GPs.
3. Theoretical Results: Appendix A.
4. Benchmark Problems: `https://hub.optuna.org/benchmarks/`.
5. Data Preprocessing: Followed `GPSampler` in Optuna 5.0. See our implementation above as well.

REFERENCES

Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A next-generation hyperparameter optimization framework. In *ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2019.

Sebastian Ament, Samuel Daulton, David Eriksson, Maximilian Balandat, and Eytan Bakshy. Unexpected improvements to expected improvement for Bayesian optimization. In *Advances in Neural Information Processing Systems*, 2023.

Maximilian Balandat, Brian Karrer, Daniel Jiang, Samuel Daulton, Ben Letham, Andrew G Wilson, and Eytan Bakshy. BoTorch: A framework for efficient Monte-Carlo Bayesian optimization. In *Advances in Neural Information Processing Systems*, 2020.

James Bergstra and Yoshua Bengio. Random search for hyper-parameter optimization. *Journal of Machine Learning Research*, 13, 2012.

James Bergstra, Brent Komer, Chris Eliasmith, Dan Yamins, and David D Cox. Hyperopt: a Python library for model selection and hyperparameter optimization. *Computational Science & Discovery*, 8, 2015.

Marek Biskup. Extrema of the two-dimensional discrete Gaussian free field. *arXiv preprint arXiv:1712.09972*, 2017.

Carlos A Coello Coello. Theoretical and numerical constraint-handling techniques used with evolutionary algorithms: a survey of the state of the art. *Computer Methods in Applied Mechanics and Engineering*, 191, 2002.

CA Coello Coello and Gregorio Toscano Pulido. Multiobjective structural optimization using a microgenetic algorithm. *Structural and Multidisciplinary Optimization*, 30, 2005.

Samuel Daulton, Maximilian Balandat, and Eytan Bakshy. Differentiable expected hypervolume improvement for parallel multi-objective Bayesian optimization. In *Advances in Neural Information Processing Systems*, 2020.

Samuel Daulton, Maximilian Balandat, and Eytan Bakshy. Parallel Bayesian optimization of multiple noisy objectives with expected hypervolume improvement. In *Advances in Neural Information Processing Systems*, 2021.

Samuel Daulton, David Eriksson, Maximilian Balandat, and Eytan Bakshy. Multi-objective Bayesian optimization over high-dimensional search spaces. In *Uncertainty in Artificial Intelligence*, 2022.

Kalyanmoy Deb and Himanshu Jain. An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, part I: Solving problems with box constraints. *IEEE Transactions on Evolutionary Computation*, 18, 2014.

Kalyanmoy Deb and Aravind Srinivasan. Innovization: Innovating design principles through optimization. In *Genetic and Evolutionary Computation Conference*, 2006.

Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and TAMT Meyarivan. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6, 2002.

Jingda Deng, Jianyong Sun, Qingfu Zhang, and Hui Li. Expected hypervolume improvement is a particular hypervolume improvement. In *AAAI Conference on Artificial Intelligence*, 2025.

Gaurav Dhiman and Vijay Kumar. Multi-objective spotted hyena optimizer: a multi-objective optimization algorithm for engineering problems. *Knowledge-Based Systems*, 150, 2018.

David Eriksson and Matthias Poloczek. Scalable constrained Bayesian optimization. In *International Conference on Artificial Intelligence and Statistics*, 2021.

Stefan Falkner, Aaron Klein, and Frank Hutter. BOHB: Robust and efficient hyperparameter optimization at scale. In *International Conference on Machine Learning*, 2018.

Azarm Farhang-Mehr and Shapour Azarm. Entropy-based multi-objective genetic algorithm for design optimization. *Structural and Multidisciplinary Optimization*, 24, 2002.

Paul Feliot, Julien Bect, and Emmanuel Vazquez. A Bayesian approach to constrained single-and multi-objective optimization. *Journal of Global Optimization*, 67, 2017.

Jacob R Gardner, Matt J Kusner, Zhixiang Eddie Xu, Kilian Q Weinberger, and John P Cunningham. Bayesian optimization with inequality constraints. In *International Conference on Machine Learning*, 2014.

Eduardo C Garrido-Merchán and Daniel Hernández-Lobato. Parallel predictive entropy search for multi-objective Bayesian optimization with constraints. *arXiv preprint arXiv:2004.00601*, 2020.

Michael A Gelbart, Jasper Snoek, and Ryan P Adams. Bayesian optimization with unknown constraints. In *Uncertainty in Artificial Intelligence*, 2014.

Daniel Hernández-Lobato, Jose Hernandez-Lobato, Amar Shah, and Ryan Adams. Predictive entropy search for multi-objective Bayesian optimization. In *International Conference on Machine Learning*, 2016.

José Miguel Hernández-Lobato, Michael Gelbart, Matthew Hoffman, Ryan Adams, and Zoubin Ghahramani. Predictive entropy search for Bayesian optimization with unknown constraints. In *International Conference on Machine Learning*, 2015.

Kaichi Irie, Shuhei Watanabe, and Masaki Onishi. Batch acquisition function evaluations and decouple optimizer updates for faster Bayesian optimization. *arXiv preprint arXiv:2511.13625*, 2025.

Himanshu Jain and Kalyanmoy Deb. An evolutionary many-objective optimization algorithm using reference-point based nondominated sorting approach, part II: Handling constraints and extending to an adaptive approach. *IEEE Transactions on Evolutionary Computation*, 18, 2013.

Jean-Pierre Kahane. Une inégalité du type de Slepian et Gordon sur les processus Gaussiens. *Israel Journal of Mathematics*, 55, 1986.

Gary R Kocis and Ignacio E Grossmann. A modelling and decomposition strategy for the MINLP optimization of process flowsheets. *Computers & Chemical Engineering*, 13, 1989.

Jing Liang, Xuanxuan Ban, Kunjie Yu, Boyang Qu, Kangjia Qiao, Caitong Yue, Ke Chen, and Kay Chen Tan. A survey on evolutionary constrained multiobjective optimization. *IEEE Transactions on Evolutionary Computation*, 27, 2022.

Jihao Andreas Lin, Sebastian Ament, Maximilian Balandat, David Eriksson, José Miguel Hernández-Lobato, and Eytan Bakshy. Scalable Gaussian processes with latent Kronecker structure. In *International Conference on Machine Learning*, 2025.

Zhi-Zhong Liu, Bing-Chuan Wang, and Ke Tang. Handling constrained multiobjective optimization problems via bidirectional coevolution. *IEEE Transactions on Cybernetics*, 52, 2022.

Sébastien Marmin, Clément Chevalier, and David Ginsbourger. Differentiating the multipoint expected improvement for optimal batch design. In *International workshop on machine learning, optimization and big data*, 2015.

Iain Murray. Differentiation of the Cholesky decomposition. *arXiv preprint arXiv:1602.07527*, 2016.

Yoshihiko Ozaki, Yuki Tanigaki, Shuhei Watanabe, and Masaki Onishi. Multiobjective tree-structured Parzen estimator for computationally expensive optimization problems. In *Genetic and Evolutionary Computation Conference*, 2020.

Yoshihiko Ozaki, Yuki Tanigaki, Shuhei Watanabe, Masahiro Nomura, and Masaki Onishi. Multiobjective tree-structured Parzen estimator. *Journal of Artificial Intelligence Research*, 73, 2022.

Yoshihiko Ozaki, Shuhei Watanabe, and Toshihiko Yanase. OptunaHub: A platform for black-box optimization. *Journal of Machine Learning Research*, 27, 2026.

Michael G Parsons and Randall L Scott. Formulation of multicriterion design optimization problems for solution with scalar numerical optimization methods. *Journal of Ship Research*, 48, 2004.

Victor Picheny. A stepwise uncertainty reduction approach to constrained global optimization. In *International Conference on Artificial Intelligence and Statistics*, 2014.

Iman Rahimi, Amir H Gandomi, Fang Chen, and Efren Mezura-Montes. A review on constraint handling techniques for population-based algorithms: from single-objective to multi-objective optimization. *Archives of Computational Methods in Engineering*, 30, 2023.

Tapabrata Ray and KM Liew. A swarm metaphor for multiobjective design optimization. *Engineering optimization*, 34, 2002.

Tapabrata Ray, Kang Tai, and Kin Chye Seow. Multiobjective design optimization by an evolutionary algorithm. *Engineering Optimization*, 33, 2001.

Ashwin Renganathan and Kade Carlson. qPOTS: Efficient batch multiobjective Bayesian optimization via Pareto optimal Thompson sampling. In *International Conference on Artificial Intelligence and Statistics*, 2025.

Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P Adams, and Nando De Freitas. Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, 104, 2015.

Nidamarthi Srinivas and Kalyanmoy Deb. Muiltiobjective optimization using nondominated sorting in genetic algorithms. *Evolutionary computation*, 2, 1994.

Masahiro Tanaka, Hikaru Watanabe, Yasuyuki Furukawa, and Tetsuzo Tanino. GA-based decision support system for multicriteria optimization. In *IEEE International Conference on Systems, Man and Cybernetics*, 1995.

Yunsheng Tian, Ane Zuniga, Xinwei Zhang, Johannes P. Dürholt, Payel Das, Jie Chen, Wojciech Matusik, and Mina Konaković Luković. Boundary exploration for Bayesian optimization with unknown physical constraints. In *International Conference on Machine Learning*, 2024.

Takashi Wada and Hideitsu Hino. Bayesian optimization for multi-objective optimization and multi-point search. *arXiv preprint arXiv:1905.02370*, 2019.

Shuhei Watanabe. Tree-structured Parzen estimator: Understanding its algorithm components and their roles for better empirical performance. *arXiv preprint arXiv:2304.11127*, 2023.

Shuhei Watanabe. Derivation of output correlation inferences for multi-output (aka multi-task) Gaussian process. *arXiv preprint arXiv:2501.07964*, 2025.

Shuhei Watanabe and Frank Hutter. c-TPE: Tree-structured Parzen estimator with inequality constraints for expensive hyperparameter optimization. In *International Joint Conference on Artificial Intelligence*, 2023.

Shuhei Watanabe and Kaichi Irie. Optuna constrained tree-structured Parzen estimator is a joint density generalization of c-TPE. *arXiv preprint arXiv:2606.09889*, 2026.

Shuhei Watanabe, Neeratyoy Mallik, Edward Bergman, and Frank Hutter. Fast benchmarking of asynchronous multi-fidelity optimization on zero-cost benchmarks. In *AutoML Conference (ABCD Track)*, 2024.

James Wilson, Frank Hutter, and Marc Deisenroth. Maximizing acquisition functions for Bayesian optimization. In *Advances in Neural Information Processing Systems*, 2018.

James T Wilson, Riccardo Moriconi, Frank Hutter, and Marc Peter Deisenroth. The reparameterization trick for acquisition functions. *arXiv preprint arXiv:1712.00424*, 2017.

Shiqiang Zhang, Robert M Lee, Behrang Shafei, David Walz, and Ruth Misener. Dependence in constrained Bayesian optimization: When do we need it and how does it help? *Optimization Letters*, 2024.

A Theory

Throughout, every probability, expectation, variance and covariance is w.r.t. the GP posterior given the history $\mathcal{D}_N$. Namely, we suppress the conditioning $\sigma(\mathcal{D}_N)$, where $\sigma(\cdot)$ is the σ-algebra generated by a map, unless it matters. Every statement is quantified over a batch $\boldsymbol{x}^{(1:Q)} \coloneqq (\boldsymbol{x}^{(1)}, \dots, \boldsymbol{x}^{(Q)}) \in \mathcal{X}^Q$, of which $\boldsymbol{x}^{(\mathcal{R})}$ is the running trials and $\boldsymbol{x}^{([Q]\setminus\mathcal{R})}$ are the candidate points, where $\mathcal{R} \coloneqq [q_{\mathrm{run}}]$ for some $q_{\mathrm{run}} \in \{0\} \cup [Q-1]$ is the running trial index set and $\mathcal{R} = \emptyset$ for $q_{\mathrm{run}} = 0$. The shuffle of $\mathcal{R}$ does not change any result below, so this canonical prefix is taken without loss of generality. The arguments are displayed only where they matter. Subset indices of $[Q]$ are written $\mathcal{I}$, and $\mathcal{I}_1, \mathcal{I}_2$ where two of them occur in the same display. The appendix is stated in terms of the *feasible observations* $Y_{N,\Omega} \coloneqq \{\boldsymbol{y}_n \in Y_N \mid \boldsymbol{c}_n \preceq \boldsymbol{0}_L\}$. Note that every variance and MSE statement is stated for S iid draws, and none of them covers the QMC draws.

A.1 Assumptions & Definitions

The following assumptions are used throughout the appendix.

Assumption 1. The joint prior over the $M+L$ latent outputs $(f_1, \dots, f_M, g_1, \dots, g_L)$ is a GP whose matrix-valued kernel k is continuously differentiable in its input arguments and strictly positive definite, and the observation noise variance satisfies $\sigma_\epsilon^2 > 0$.

Assumption 2. The reference point $\boldsymbol{r} \in \mathbb{R}^M$ is finite, and $Y_{N,\Omega}$ is a finite subset of $\mathcal{Y} \coloneqq (-\infty, \boldsymbol{r})$.

For notational simplicity, we use the following definitions.

Definition 4 (Hypervolume Improvement)**.** Let the $\mathcal{H}^{(\emptyset)}$-measurable *hypervolume improvement* of the batch points indexed by any $\mathcal{I} \subseteq [Q]$ be

$$\delta V^{(\mathcal{I})} \coloneqq \mathrm{vol}\Big(\bigcap_{q\in\mathcal{I}} [\boldsymbol{f}^{(q)}, \boldsymbol{r}) \setminus \bigcup_{\boldsymbol{y}\in Y_{N,\Omega}} [\boldsymbol{y}, \boldsymbol{r})\Big) \geq 0, \tag{9}$$

For $\mathcal{I} \neq \emptyset$ and any $q \in \mathcal{I}$, Eq. (9) bounds $0 \leq \delta V^{(\mathcal{I})} \leq \prod_{m=1}^{M} \mathtt{relu}(r_m - f_m^{(q)})$, and Eq. (12)-Eq. (13) below bound $\Delta^{(Q)}$ by the same expression at $q = Q$ and B by the sum of those expressions over $q \in [Q-1]$. Each factor is finite by Assumption 2 and has moments of every order under the Gaussian posterior, so generalized Hölder's inequality places $\delta V^{(\mathcal{I})}$, $\Delta^{(Q)}$, $B \in L^p$ for every $p \geq 1$, and every expectation, variance and covariance in this appendix is finite.

Definition 5 (The Rao-Blackwell Conditioning Lattice)**.** Define

$$\mathcal{H}^{(\emptyset)} \coloneqq \sigma\big(\boldsymbol{f}^{(1:Q)}\big), \qquad \mathcal{H}^{(\mathcal{R})} \coloneqq \sigma\big(\boldsymbol{f}^{(1:Q)}, \boldsymbol{g}^{(\mathcal{R})}\big), \qquad \mathcal{H}_g \coloneqq \sigma\big(\boldsymbol{g}^{(1:Q-1)}\big). \tag{10}$$

The *Rao-Blackwell conditioning lattice* is the family of σ-algebras $\mathcal{H}$ with $\mathcal{H}^{(\emptyset)} \subseteq \mathcal{H} \subseteq \mathcal{H}^{(1:Q-1)}$. Its top $\mathcal{H}^{(1:Q-1)}$ is everything that is sampled except the candidate constraint values $\boldsymbol{g}^{(Q)}$. Its bottom $\mathcal{H}^{(\emptyset)}$ retains the objective fantasy samples alone and marginalizes every constraint latent.

For nested $\mathcal{H}_1 \subseteq \mathcal{H}_2$, an integrable X_1 and a $\mathcal{H}_1$-measurable X_2 with X_1X_2 integrable, we use the tower property $\mathbb{E}[\mathbb{E}[X_1 \mid \mathcal{H}_2] \mid \mathcal{H}_1] = \mathbb{E}[X_1 \mid \mathcal{H}_1]$ and the pull-out property $\mathbb{E}[X_1X_2 \mid \mathcal{H}_1] = X_2\mathbb{E}[X_1 \mid \mathcal{H}_1]$ throughout, both a.s.

Definition 6 (Fantasy Pareto Set & Candidate Hypervolume Improvement & Baseline)**.**

$$Y_{N,\Omega}^{(1:q)} \coloneqq Y_{N,\Omega} \cup \big\{\boldsymbol{f}^{(q')} \,\big|\, \boldsymbol{g}^{(q')} \preceq \mathbf{0}_L,\ q' \in [q]\big\}, \tag{11}$$

$$\Delta^{(Q)} \coloneqq \mathrm{HVI}^{(Q)} = V\big(Y_{N,\Omega}^{(1:Q-1)} \cup \{\boldsymbol{f}^{(Q)}\}, \boldsymbol{r}\big) - V\big(Y_{N,\Omega}^{(1:Q-1)}, \boldsymbol{r}\big) \geq 0, \tag{12}$$

$$B \coloneqq V\big(Y_{N,\Omega}^{(1:Q-1)}, \boldsymbol{r}\big) - V(Y_{N,\Omega}, \boldsymbol{r}) \geq 0. \tag{13}$$

Both are $\mathcal{H}^{(1:Q-1)}$-measurable, and we write $Y_{N,\Omega,s}^{(1:Q-1)}, \Delta_s^{(Q)}, B_s$ for the s-th fantasy sample.

Define $\mathbb{I}_\Omega^{(\mathcal{I})} \coloneqq \prod_{q\in\mathcal{I}} \mathbb{I}_\Omega[\boldsymbol{g}^{(q)} \preceq \mathbf{0}_L]$ so that $\mathbb{I}_\Omega^{(\mathcal{I}_1)}\mathbb{I}_\Omega^{(\mathcal{I}_2)} = \mathbb{I}_\Omega^{(\mathcal{I}_1\cup\mathcal{I}_2)}$ for disjoint $\mathcal{I}_1, \mathcal{I}_2$. The conditional PoF is written as $p_\Omega^{(\mathcal{I})}(\mathcal{H}) \coloneqq \mathbb{P}[\mathbb{I}_\Omega^{(\mathcal{I})} \mid \mathcal{H}]$ given a lattice member $\mathcal{H}$. Write $p_\Omega^{(Q)}(\mathcal{H}) \coloneqq p_\Omega^{(\{Q\})}(\mathcal{H})$ and $\mathbb{I}_\Omega^{(Q)} \coloneqq \mathbb{I}_\Omega^{(\{Q\})}$ for the singleton case. At $\mathcal{H}^{(\emptyset)}$, $\mathcal{H}^{(1:Q-1)}$ and $\mathcal{H}_g$, the conditional law of the remaining fantasy samples is Gaussian, and $p_\Omega^{(\mathcal{I})}(\mathcal{H})$ denotes that given by its $|\mathcal{I}|L$-dimensional Gaussian CDF.

Lemma 7 (Hard Indicator Utility u_{hard})**.** *For every $\boldsymbol{x}^{(1:Q)} \in \mathcal{X}^Q$ and pathwise,*

$$V(Y_{N,\Omega}^{(1:Q)}, \boldsymbol{r}) - V(Y_{N,\Omega}, \boldsymbol{r}) = \sum_{\emptyset\neq\mathcal{I}\subseteq[Q]} (-1)^{|\mathcal{I}|+1} \delta V^{(\mathcal{I})} \mathbb{I}_\Omega^{(\mathcal{I})} = B + \Delta^{(Q)}\mathbb{I}_\Omega^{(Q)} = u_{\mathrm{hard}}. \tag{14}$$

Proof. Let $R \coloneqq \bigcup_{\boldsymbol{y}\in Y_{N,\Omega}} [\boldsymbol{y}, \boldsymbol{r})$ be the region dominated by the feasible observations and let

$$A_q \coloneqq [\boldsymbol{f}^{(q)}, \boldsymbol{r}) \setminus R\,(q \in [Q]), \text{ so that } \mathrm{vol}\Big(\bigcap_{q\in\mathcal{I}} A_q\Big) = \delta V^{(\mathcal{I})} \text{ for } \mathcal{I} \neq \emptyset \tag{15}$$

by Eq. (9), every such volume being finite because $\boldsymbol{r}$ is finite. Let $\mathcal{I}_\Omega \coloneqq \{q \in [Q] \mid \boldsymbol{g}^{(q)} \preceq \mathbf{0}_L\}$ be the realized feasible index set, for which $Y_{N,\Omega}^{(1:Q)} = Y_{N,\Omega} \cup \{\boldsymbol{f}^{(q)}\}_{q\in\mathcal{I}_\Omega}$ by Eq. (11) and $\mathbb{I}_\Omega^{(\mathcal{I})} = \mathbb{I}_\Omega[\mathcal{I} \subseteq \mathcal{I}_\Omega]$. Since $V(A, \boldsymbol{r}) = \mathrm{vol}(\bigcup_{\boldsymbol{y}\in A}[\boldsymbol{y}, \boldsymbol{r}))$ and $\bigcup_{\boldsymbol{y}\in Y_{N,\Omega}^{(1:Q)}} [\boldsymbol{y}, \boldsymbol{r}) = R \sqcup \bigcup_{q\in\mathcal{I}_\Omega} A_q$, inclusion-exclusion on that finite union gives

$$\begin{aligned} V\big(Y_{N,\Omega}^{(1:Q)}, \boldsymbol{r}\big) - V(Y_{N,\Omega}, \boldsymbol{r}) &= \mathrm{vol}\Big(\bigcup_{q\in\mathcal{I}_\Omega} A_q\Big) \\ &= \sum_{\emptyset\neq\mathcal{I}\subseteq\mathcal{I}_\Omega} (-1)^{|\mathcal{I}|+1}\delta V^{(\mathcal{I})} = \sum_{\emptyset\neq\mathcal{I}\subseteq[Q]} (-1)^{|\mathcal{I}|+1}\delta V^{(\mathcal{I})}\mathbb{I}_\Omega^{(\mathcal{I})}, \end{aligned} \tag{16}$$

the last step because $\mathbb{I}_\Omega^{(\mathcal{I})}$ annihilates exactly the terms whose $\mathcal{I}$ contains an infeasible index. This is the first equality of Eq. (14), i.e., the $2^Q - 1$ subset enumeration of Daulton et al. (2020) written pathwise.

Now split that sum at $Q \in \mathcal{I}$. The terms with $Q \notin \mathcal{I}$ are Eq. (16) run over the sub-batch $\boldsymbol{x}^{(1:Q-1)}$, whose realized feasible index set is $\mathcal{I}_\Omega \cap [Q-1]$ and whose augmented set is $Y_{N,\Omega}^{(1:Q-1)}$, so they equal B by Eq. (13). The terms with $Q \in \mathcal{I}$ are therefore $V(Y_{N,\Omega}^{(1:Q)}, \boldsymbol{r}) - V(Y_{N,\Omega}^{(1:Q-1)}, \boldsymbol{r})$, which is $\Delta^{(Q)}$ on $\{\mathbb{I}_\Omega^{(Q)} = 1\}$, and 0 on $\{\mathbb{I}_\Omega^{(Q)} = 0\}$. This is $\Delta^{(Q)}\mathbb{I}_\Omega^{(Q)}$ by Eq. (12). □

Definition 8 (Utilities and AFs)**.** For $\bullet \in \{\mathrm{RB}, \mathrm{marg}, \mathrm{hard}, \mathrm{sig}\}$, the utilities and AFs of interest are represented as:

$$u_\bullet(\mathcal{H}) \coloneqq \sum_{\emptyset \neq \mathcal{I} \subseteq [Q]} (-1)^{|\mathcal{I}|+1}\, \delta V^{(\mathcal{I})}\, \pi_\bullet^{(\mathcal{I})}(\mathcal{H}), \quad \alpha_\bullet(\mathcal{H}) \coloneqq \mathbb{E}[u_\bullet(\mathcal{H})], \quad \hat{\alpha}_{\bullet,S}(\mathcal{H}) \coloneqq \frac{1}{S}\sum_{s=1}^{S} u_{\bullet,s}(\mathcal{H}), \tag{17}$$

where $u_{\bullet,s}$ is $u_\bullet$ at the s-th reparametrized fantasy draw under the S-sample Monte Carlo estimation and each utility uses the following PoF representations:

$$\pi_{\mathrm{RB}}^{(\mathcal{I})} = p_\Omega^{(\mathcal{I})}(\mathcal{H}),\ \pi_{\mathrm{marg}}^{(\mathcal{I})} = \mathbb{I}_\Omega^{(\mathcal{I}\cap\mathcal{R})} p_\Omega^{(\mathcal{I}\setminus\mathcal{R})},\ \pi_{\mathrm{hard}}^{(\mathcal{I})} = \mathbb{I}_\Omega^{(\mathcal{I})},\ \pi_{\mathrm{sig}}^{(\mathcal{I})} = \mathbb{I}_\Omega^{(\mathcal{I}\cap\mathcal{R})} \prod_{q\in\mathcal{I}\setminus\mathcal{R}} \prod_{l=1}^{L} \varsigma\big(-g_l^{(q)}/\tau\big). \tag{18}$$

The case $\mathcal{R} = [Q-1]$, for which $\mathcal{H}^{(\mathcal{R})} = \mathcal{H}^{(1:Q-1)}$, is of primary interest to us. Using Lemma 7, the utilities then reduce to:

$$u_\bullet(\mathcal{H}^{(1:Q-1)}) = B + \Delta^{(Q)} \pi_\bullet^{(Q)}(\mathcal{H}^{(1:Q-1)}) \tag{19}$$

Note $u_{\mathrm{RB}}(\mathcal{H})$ needs only the fantasy samples that build $\mathcal{H}$, so $u_{\mathrm{RB}}(\mathcal{H}^{(1:Q-1)})$ needs $(\boldsymbol{f}^{(1:Q)}, \boldsymbol{g}^{(1:Q-1)})$ and $u_{\mathrm{RB}}(\mathcal{H}^{(\emptyset)})$ needs $\boldsymbol{f}^{(1:Q)}$ alone, whereas u_{hard} and u_{sig} also need $\boldsymbol{g}^{(Q)}$.

A.2 Exactness of Conditional PoF by Rao-Blackwellization

Theorem 9 (Exact Conditional PoF by Rao-Blackwellization)**.** *For every $\boldsymbol{x}^{(1:Q)} \in \mathcal{X}^Q$ and $\mathcal{H}$ s.t. $\mathcal{H}^{(\emptyset)} \subseteq \mathcal{H} \subseteq \mathcal{H}^{(1:Q-1)}$:*

(i) ***Exact Conditional Expectation****: $\mathbb{E}[u_{\mathrm{hard}} \mid \mathcal{H}] = u_{\mathrm{RB}}(\mathcal{H})$ a.s.*

(ii) ***Exact Identity****: $\alpha_{\mathrm{RB}}(\mathcal{H}) = \alpha_{\mathrm{hard}}$ as an identity of functions on $\mathcal{X}^Q$.*

(iii) ***Nesting****: for $\mathcal{H}^{(\emptyset)} \subseteq \mathcal{H}_1 \subseteq \mathcal{H}_2 \subseteq \mathcal{H}^{(1:Q-1)}$, $u_{\mathrm{RB}}(\mathcal{H}_1) = \mathbb{E}[u_{\mathrm{RB}}(\mathcal{H}_2) \mid \mathcal{H}_1]$ a.s.*

Proof. (i) By Eq. (9), each $\delta V^{(\mathcal{I})}$ is a Borel function of $\boldsymbol{f}^{(1:Q)}$. Thus, it is $\mathcal{H}^{(\emptyset)}$-measurable and $\mathcal{H}$-measurable for every lattice member, and integrable by Assumption 2. Since $\mathbb{I}_\Omega^{(\mathcal{I})} \in \{0,1\}$, the product is integrable too and the pull-out property gives $\mathbb{E}[\delta V^{(\mathcal{I})}\mathbb{I}_\Omega^{(\mathcal{I})} \mid \mathcal{H}] = \delta V^{(\mathcal{I})} p_\Omega^{(\mathcal{I})}(\mathcal{H})$ a.s. Applying it to each term of Eq. (14) and summing by linearity yields $\mathbb{E}[u_{\mathrm{hard}} \mid \mathcal{H}] = u_{\mathrm{RB}}(\mathcal{H})$ a.s.

(ii) Taking expectations in (i) and using the tower property, $\alpha_{\mathrm{hard}} = \mathbb{E}[u_{\mathrm{hard}}] = \mathbb{E}[\mathbb{E}[u_{\mathrm{hard}} \mid \mathcal{H}]] = \mathbb{E}[u_{\mathrm{RB}}(\mathcal{H})] = \alpha_{\mathrm{RB}}(\mathcal{H})$, at each fixed $\boldsymbol{x}^{(1:Q)} \in \mathcal{X}^Q$ separately.

(iii) By (i) at $\mathcal{H}_1$ and at $\mathcal{H}_2$, both members of the lattice, and the tower property, $u_{\mathrm{RB}}(\mathcal{H}_1) = \mathbb{E}[u_{\mathrm{hard}} \mid \mathcal{H}_1] = \mathbb{E}[\mathbb{E}[u_{\mathrm{hard}} \mid \mathcal{H}_2] \mid \mathcal{H}_1] = \mathbb{E}[u_{\mathrm{RB}}(\mathcal{H}_2) \mid \mathcal{H}_1]$ a.s. □

A.3 The Exact Variance Gap along the Conditioning Lattice

Theorem 10 (Variance Reduction with Exact Gap)**.** *For every $\boldsymbol{x}^{(1:Q)} \in \mathcal{X}^Q$ and $\mathcal{H}$ s.t. $\mathcal{H}^{(\emptyset)} \subseteq \mathcal{H} \subseteq \mathcal{H}^{(1:Q-1)}$:*

(i) ***Exact Gap****:*

$$\mathbb{V}[u_{\mathrm{hard}}] - \mathbb{V}[u_{\mathrm{RB}}(\mathcal{H})] = \mathbb{E}\big[\mathbb{V}[u_{\mathrm{hard}} \mid \mathcal{H}]\big] \geq 0, \tag{20}$$

$$\mathbb{V}[u_{\mathrm{hard}} \mid \mathcal{H}] = \sum_{\emptyset \neq \mathcal{I}_1, \mathcal{I}_2 \subseteq [Q]} (-1)^{|\mathcal{I}_1|+|\mathcal{I}_2|}\, \delta V^{(\mathcal{I}_1)} \delta V^{(\mathcal{I}_2)} \big(p_\Omega^{(\mathcal{I}_1 \cup \mathcal{I}_2)}(\mathcal{H}) - p_\Omega^{(\mathcal{I}_1)}(\mathcal{H})\, p_\Omega^{(\mathcal{I}_2)}(\mathcal{H})\big). \tag{21}$$

(ii) ***Lattice Monotonicity****: for $\mathcal{H}^{(\emptyset)} \subseteq \mathcal{H}_1 \subseteq \mathcal{H}_2 \subseteq \mathcal{H}^{(1:Q-1)}$,*

$$\mathbb{V}[u_{\mathrm{RB}}(\mathcal{H}^{(\emptyset)})] \leq \mathbb{V}[u_{\mathrm{RB}}(\mathcal{H}_1)] \leq \mathbb{V}[u_{\mathrm{RB}}(\mathcal{H}_2)] \leq \mathbb{V}[u_{\mathrm{hard}}], \tag{22}$$

where each is given exactly by $\mathbb{V}[u_{\mathrm{RB}}(\mathcal{H}_1)] = \mathbb{V}[u_{\mathrm{RB}}(\mathcal{H}_2)] - \mathbb{E}[\mathbb{V}[u_{\mathrm{RB}}(\mathcal{H}_2) \mid \mathcal{H}_1]]$.

Proof. Every variance below is finite because u_{hard} and $u_{\mathrm{RB}}(\mathcal{H})$ are finite subset sums of products of an L^2 factor $\delta V^{(\mathcal{I})}$ (Assumption 2) and a factor in $[0,1]$.

(i) Since $\mathbb{E}[u_{\mathrm{hard}} \mid \mathcal{H}] = u_{\mathrm{RB}}(\mathcal{H})$ a.s. by Theorem 9 (i), the law of total variance reads $\mathbb{V}[u_{\mathrm{hard}}] = \mathbb{V}[u_{\mathrm{RB}}(\mathcal{H})] + \mathbb{E}[\mathbb{V}[u_{\mathrm{hard}} \mid \mathcal{H}]]$, and rearranging gives Eq. (20). Expanding u_{hard} by Eq. (14) inside $\mathbb{V}[\cdot \mid \mathcal{H}]$ and pulling the $\mathcal{H}$-measurable $\delta V^{(\mathcal{I})}$ out leaves $\mathrm{Cov}[\mathbb{I}_\Omega^{(\mathcal{I}_1)}, \mathbb{I}_\Omega^{(\mathcal{I}_2)} \mid \mathcal{H}]$ in each term.

(ii) By Theorem 9 (iii) and the same law of total variance, now at $\mathcal{H}_1$, the stated gap is $\mathbb{E}[\mathbb{V}[u_{\mathrm{RB}}(\mathcal{H}_2) \mid \mathcal{H}_1]] \geq 0$, which is Eq. (22) at $(\mathcal{H}^{(\emptyset)}, \mathcal{H}_1)$ and at $(\mathcal{H}_1, \mathcal{H}_2)$, its last inequality being (i) at $\mathcal{H} = \mathcal{H}_2$. □

Corollary 11 (IID-MC Estimator Variance)**.** *For S* iid *fantasy samples and any $\mathcal{H}$ of the lattice,* $\mathbb{V}[\hat{\alpha}_{\mathrm{hard},S}] - \mathbb{V}[\hat{\alpha}_{\mathrm{RB},S}(\mathcal{H})] = S^{-1}\,\mathbb{E}[\mathbb{V}[u_{\mathrm{hard}} \mid \mathcal{H}]] \geq 0$.

Proof. The draws being iid, $\mathbb{V}[\hat{\alpha}_{\bullet,S}] = S^{-1}\mathbb{V}[u_\bullet]$. Apply Eq. (20) to $\bullet \in \{\mathrm{hard}, \mathrm{RB}\}$. □

A.4 The Exact Bias of the Marginal-PoF Utility

The following theorem gives the utility error in closed form at $Q > 1$ when using the marginal PoF instead of the feasibility indicator.

Theorem 12 (The Exact Bias of the Marginal-PoF Utility)**.** *For every $\boldsymbol{x}^{(1:Q)} \in \mathcal{X}^Q$ and $\mathcal{H}^{(\emptyset)} \subseteq \mathcal{H}^{(\mathcal{R})} \subseteq \mathcal{H}^{(1:Q-1)}$, the AF values before the log-transformation satisfy*

$$\alpha_{\mathrm{marg}}(\mathcal{H}^{(\mathcal{R})}) - \alpha_{\mathrm{hard}} = -\sum_{\emptyset \neq \mathcal{I} \subseteq [Q]} (-1)^{|\mathcal{I}|+1}\,\mathrm{Cov}\big[\delta V^{(\mathcal{I})}\mathbb{I}_\Omega^{(\mathcal{I}\cap\mathcal{R})},\ \mathbb{I}_\Omega^{(\mathcal{I}\setminus\mathcal{R})}\big], \tag{23}$$

$$\mathrm{Cov}\big[\delta V^{(\mathcal{I})}\mathbb{I}_\Omega^{(\mathcal{I}\cap\mathcal{R})},\ \mathbb{I}_\Omega^{(\mathcal{I}\setminus\mathcal{R})}\big] = \mathrm{Cov}\big[\delta V^{(\mathcal{I})}\mathbb{I}_\Omega^{(\mathcal{I}\cap\mathcal{R})},\ p_\Omega^{(\mathcal{I}\setminus\mathcal{R})}(\mathcal{H}^{(\mathcal{R})})\big].$$

Proof. $\alpha_{\mathrm{marg}}(\mathcal{H}^{(\mathcal{R})}) = \sum_{\emptyset\neq\mathcal{I}\subseteq[Q]}(-1)^{|\mathcal{I}|+1}\mathbb{E}[\delta V^{(\mathcal{I})}\mathbb{I}_\Omega^{(\mathcal{I}\cap\mathcal{R})}]\,\mathbb{E}[\mathbb{I}_\Omega^{(\mathcal{I}\setminus\mathcal{R})}]$. The product form of $\mathbb{I}_\Omega^{(\mathcal{I})}$ gives $\mathbb{I}_\Omega^{(\mathcal{I}\cap\mathcal{R})}\mathbb{I}_\Omega^{(\mathcal{I}\setminus\mathcal{R})} = \mathbb{I}_\Omega^{(\mathcal{I})}$, so $\alpha_{\mathrm{hard}} = \sum_{\emptyset\neq\mathcal{I}\subseteq[Q]}(-1)^{|\mathcal{I}|+1}\mathbb{E}[\delta V^{(\mathcal{I})}\mathbb{I}_\Omega^{(\mathcal{I}\cap\mathcal{R})}\mathbb{I}_\Omega^{(\mathcal{I}\setminus\mathcal{R})}]$. Subtracting termwise, $\mathbb{E}[A]\mathbb{E}[B] - \mathbb{E}[AB] = -\mathrm{Cov}[A,B]$ at $A \coloneqq \delta V^{(\mathcal{I})}\mathbb{I}_\Omega^{(\mathcal{I}\cap\mathcal{R})}$ and $B \coloneqq \mathbb{I}_\Omega^{(\mathcal{I}\setminus\mathcal{R})}$ is the first identity. For the second, A is $\mathcal{H}^{(\mathcal{R})}$-measurable and integrable, so the pull-out and tower properties give both $\mathbb{E}[AB] = \mathbb{E}[A\,p_\Omega^{(\mathcal{I}\setminus\mathcal{R})}(\mathcal{H}^{(\mathcal{R})})]$ and $\mathbb{E}[B] = \mathbb{E}[p_\Omega^{(\mathcal{I}\setminus\mathcal{R})}(\mathcal{H}^{(\mathcal{R})})]$. This completes the proof. □

Remark 13 (Sign and argmax Consequences). u_{marg} overestimates where $\Delta^{(Q)}$ and the conditional PoF are negatively associated and underestimates where they are positively associated. The covariance depends on $\boldsymbol{x}^{(1:Q)}$, so the bias need not be a constant offset of α_{hard} and $\mathrm{argmax}\,\alpha_{\mathrm{marg}} = \mathrm{argmax}\,\alpha_{\mathrm{hard}}$ is not guaranteed. The operative failure mode for BO is not a tolerably wrong value but the wrong candidate.

Besides $Q = 1$, the following is a sufficient condition for the marginal PoF utility to be exact.

Proposition 14 (When the Marginal-PoF Utility Is Exact)**.** *Fix $\boldsymbol{x}^{(1:Q)} \in \mathcal{X}^Q$ and let $\mathcal{R} = [Q-1]$. If $\mathrm{Cov}\big[\boldsymbol{g}^{(Q)}, (\boldsymbol{f}^{(1:Q)}, \boldsymbol{g}^{(1:Q-1)})\big] = \mathbf{0}$, then $\alpha_{\mathrm{marg}} = \alpha_{\mathrm{hard}}$.*

Proof. Write $p_z \coloneqq p_\Omega^{(Q)}(\mathcal{H}^{(1:Q-1)})$ and $p_N \coloneqq p_\Omega^{(Q)}$, the latter deterministic. Applying the argument of Theorem 12 to $(\Delta^{(Q)}, \mathbb{I}_\Omega^{(Q)})$ gives the following:

$$\alpha_{\mathrm{marg}} - \alpha_{\mathrm{hard}} = \mathbb{E}\big[\Delta^{(Q)}\big]\,p_N - \mathbb{E}\big[\Delta^{(Q)}\,\mathbb{I}_\Omega^{(Q)}\big] = -\mathrm{Cov}\big[\Delta^{(Q)},\ \mathbb{I}_\Omega^{(Q)}\big] = -\mathrm{Cov}\big[\Delta^{(Q)},\ p_z\big]. \tag{24}$$

It therefore suffices that p_z be constant a.s. The latents are jointly Gaussian, so the vanishing cross-covariance is independence of $\boldsymbol{g}^{(Q)}$ from $(\boldsymbol{f}^{(1:Q)}, \boldsymbol{g}^{(1:Q-1)})$. Therefore, $p_z = p_N$ a.s. □

Remark 15 (An Estimate of the Bias at No Extra Sampling). Both factors of the covariance in Eq. (24) are $\mathcal{H}^{(1:Q-1)}$-measurable, so the pairs $(\Delta_s^{(Q)}, p_s)$ determine an empirical estimate of $\alpha_{\mathrm{marg}} - \alpha_{\mathrm{hard}}$ by $\frac{-1}{S-1}\sum_{s=1}^{S}(\Delta_s^{(Q)} - \overline{\Delta^{(Q)}})(p_s - \bar{p})$ where $\bar{\cdot}$ represents the empirical mean over the fantasy samples and p_s is $p_\Omega^{(Q)}(\mathcal{H}^{(1:Q-1)})$ given the s-th fantasy sample. The estimate is unbiased for iid fantasy samples and reads off the fantasy samples for $\hat{\alpha}_{\mathrm{RB},S}(\mathcal{H}^{(1:Q-1)})$, so it costs no extra sampling. A positive estimate diagnoses $\alpha_{\mathrm{marg}} > \alpha_{\mathrm{hard}}$, i.e., the overestimation in Remark 20. Let $p_N \coloneqq p_\Omega^{(Q)}$. Combining Eq. (24) with the Cauchy-Schwarz inequality and $\mathbb{V}[\mathbb{I}_\Omega^{(Q)}] = p_N(1-p_N)$ gives $|\alpha_{\mathrm{marg}} - \alpha_{\mathrm{hard}}| \leq \sqrt{\mathbb{V}[\Delta^{(Q)}]\,p_N(1-p_N)}$. At a fixed $\mathbb{V}[\Delta^{(Q)}]$, this bound vanishes as $p_N \to \{0,1\}$ and peaks at $p_N = 1/2$, so an uncertain marginal PoF is where the identity permits the largest bias.

Since the bias of Theorem 12 does not depend on how many draws are spent, it survives in the mean squared error against the exact target.

Corollary 16 (Mean Squared Error Against the Exact Target α_{hard})**.** *For S iid draws, every $\boldsymbol{x}^{(1:Q)} \in \mathcal{X}^Q$ and $\mathcal{H}^{(\emptyset)} \subseteq \mathcal{H}^{(\mathcal{R})} \subseteq \mathcal{H}^{(1:Q-1)}$, let both estimators be scored against the same exact target by $\mathrm{MSE}[\hat{\alpha}_{\bullet,S}] \coloneqq \mathbb{E}\big[(\hat{\alpha}_{\bullet,S} - \alpha_{\mathrm{hard}})^2\big]$:*

(i) ***Bias Floor of the Marginal-PoF Utility****:*

$$\begin{aligned}\mathrm{MSE}\big[\hat{\alpha}_{\mathrm{marg},S}(\mathcal{H}^{(\mathcal{R})})\big] &= \big(\alpha_{\mathrm{marg}}(\mathcal{H}^{(\mathcal{R})}) - \alpha_{\mathrm{hard}}\big)^2 + S^{-1}\,\mathbb{V}[u_{\mathrm{marg}}(\mathcal{H}^{(\mathcal{R})})] \\ &\geq \big(\alpha_{\mathrm{marg}}(\mathcal{H}^{(\mathcal{R})}) - \alpha_{\mathrm{hard}}\big)^2,\end{aligned} \tag{25}$$

a lower bound independent of S.

(ii) ***Vanishing Error of the Rao-Blackwellized Estimator****:*

$$\mathrm{MSE}[\hat{\alpha}_{\mathrm{RB},S}(\mathcal{H}^{(\mathcal{R})})] = S^{-1}\,\mathbb{V}[u_{\mathrm{RB}}(\mathcal{H}^{(\mathcal{R})})] \to 0 \quad (S \to \infty). \tag{26}$$

(iii) ***Dominance for Sufficiently Many Samples****: if $\alpha_{\mathrm{marg}}(\mathcal{H}^{(\mathcal{R})}) \neq \alpha_{\mathrm{hard}}$,*

$$\mathrm{MSE}[\hat{\alpha}_{\mathrm{RB},S}(\mathcal{H}^{(\mathcal{R})})] < \mathrm{MSE}[\hat{\alpha}_{\mathrm{marg},S}(\mathcal{H}^{(\mathcal{R})})], \tag{27}$$

for every $S > \mathbb{V}[u_{\mathrm{RB}}(\mathcal{H}^{(\mathcal{R})})]/(\alpha_{\mathrm{marg}}(\mathcal{H}^{(\mathcal{R})}) - \alpha_{\mathrm{hard}})^2$.

Proof. Every variance below is finite by the moment bound stated after Definition 4 under Assumption 2. The draws being iid, $\mathbb{E}[\hat{\alpha}_{\bullet,S}] = \alpha_\bullet$ and $\mathbb{V}[\hat{\alpha}_{\bullet,S}] = S^{-1}\mathbb{V}[u_\bullet]$, as in the proof of Corollary 11, so the bias-variance decomposition about α_{hard} reads

$$\mathrm{MSE}\big[\hat{\alpha}_{\bullet,S}\big] = \big(\alpha_\bullet - \alpha_{\mathrm{hard}}\big)^2 + S^{-1}\,\mathbb{V}[u_\bullet]. \tag{28}$$

(i) is Eq. (28) at $\bullet = \mathrm{marg}(\mathcal{H}^{(\mathcal{R})})$, whose squared-bias term is the square of the covariance sum Eq. (23) of Theorem 12, dropped to the stated bound by $\mathbb{V}[u_{\mathrm{marg}}(\mathcal{H}^{(\mathcal{R})})] \geq 0$ alone.

(ii) is Eq. (28) at $\bullet = \mathrm{RB}(\mathcal{H}^{(\mathcal{R})})$, whose bias term vanishes because $\alpha_{\mathrm{RB}}(\mathcal{H}^{(\mathcal{R})}) = \alpha_{\mathrm{hard}}$ by Theorem 9 (ii).

(iii) By (ii) and the bound of (i), $\mathrm{MSE}[\hat{\alpha}_{\mathrm{RB},S}(\mathcal{H}^{(\mathcal{R})})] < \mathrm{MSE}[\hat{\alpha}_{\mathrm{marg},S}(\mathcal{H}^{(\mathcal{R})})]$ holds as soon as $S^{-1}\mathbb{V}[u_{\mathrm{RB}}(\mathcal{H}^{(\mathcal{R})})] < (\alpha_{\mathrm{marg}}(\mathcal{H}^{(\mathcal{R})}) - \alpha_{\mathrm{hard}})^2$, which is the stated threshold on S, the right-hand side being nonzero by hypothesis. Both estimators are formed from the same S draws of $(\boldsymbol{f}^{(1:Q)}, \boldsymbol{g}^{(\mathcal{R})})$, so the comparison is at matched sampling cost. □

A.5 Pathwise Gradients of the Rao-Blackwellized Utility

Theorem 9 (ii) is an identity between expectations, equating two AF values at every $\boldsymbol{x}^{(1:Q)}$. Gradient-based methods consume the pathwise gradient $\nabla u_{\bullet,s}$ of the realized utility on each fantasy sample. On such a draw, every $\mathbb{I}_\Omega^{(\mathcal{I})}$ is piecewise constant in $\boldsymbol{x}^{(1:Q)}$, so Eq. (14) gives $\nabla u_{\mathrm{hard}} = \sum_{\mathcal{I}} (-1)^{|\mathcal{I}|+1} \mathbb{I}_\Omega^{(\mathcal{I})} \nabla \delta V^{(\mathcal{I})} = \mathbb{I}_\Omega^{(Q)} \nabla \Delta^{(Q)}$ w.r.t. $\boldsymbol{x}^{(Q)}$, and the candidate-block pathwise gradient vanishes whenever the sampled candidate is infeasible. The following theorem establishes the corresponding result for every $\mathcal{H}^{(\mathcal{R})}$ in Eq. (10) with $q_{\mathrm{run}} \in \{0\} \cup [Q-1]$. That is, it applies specifically to the Gaussian-coordinate sublattice. For notational simplicity, define the candidate index set as $\mathcal{C} \coloneqq [Q] \setminus \mathcal{R}$ in this section. Additionally, we define $\mathcal{I}_\mathcal{C} = \mathcal{I} \cap \mathcal{C} = \mathcal{I} \setminus \mathcal{R}$ and $\mathcal{I}_\mathcal{R} = \mathcal{I} \cap \mathcal{R}$.

Theorem 17 (Pathwise Gradient of the the RB Approach's Utility)**.** *Fix fantasy samples, and assume k admits a C^1 extension to an open neighborhood of $\mathcal{X} \times \mathcal{X} \subseteq \mathbb{R}^D \times \mathbb{R}^D$. Let the fixed inputs $\boldsymbol{x}^{(\mathcal{R})}$ be pairwise distinct and held fixed during AF optimization. Then on the open set $\{\boldsymbol{x}^{(\mathcal{C})} \in \mathcal{X}^{Q-q_{\mathrm{run}}} \mid \boldsymbol{x}^{(1:Q)} \in \mathcal{X}^Q_{\neq}\}$:*

(i) Conditional-PoF Factorization*: for every $\emptyset \neq \mathcal{I} \subseteq [Q]$, $p_\Omega^{(\mathcal{I})}(\mathcal{H}^{(\mathcal{R})}) = \mathbb{I}_\Omega^{(\mathcal{I}_\mathcal{R})}\, p_\Omega^{(\mathcal{I}_\mathcal{C})}(\mathcal{H}^{(\mathcal{R})})$ a.s., and for $\mathcal{I} \not\subseteq \mathcal{R}$ the conditional law of $(\boldsymbol{g}^{(q)})_{q \in \mathcal{I}_\mathcal{C}}$ given $\mathcal{H}^{(\mathcal{R})}$ is a non-degenerate $|\mathcal{I}_\mathcal{C}|$-L-variate Gaussian whose mean and covariance are C^1 in $\boldsymbol{x}^{(\mathcal{C})}$, so that $p_\Omega^{(\mathcal{I}_\mathcal{C})}(\mathcal{H}^{(\mathcal{R})})$ is C^1 there with values in $(0,1)$.*

(ii) Gradient*: every realized $\mathbb{I}_\Omega^{(\mathcal{I}_\mathcal{R})}$ is constant in $\boldsymbol{x}^{(\mathcal{C})}$, and the map $\boldsymbol{x}^{(\mathcal{C})} \mapsto u_{\mathrm{RB}}(\mathcal{H}^{(\mathcal{R})})$ is locally Lipschitz, differentiable at Lebesgue-a.e. $\boldsymbol{x}^{(\mathcal{C})}$, with*

$$\nabla_{\boldsymbol{x}^{(\mathcal{C})}} u_{\mathrm{RB}}(\mathcal{H}^{(\mathcal{R})}) = \sum_{\emptyset \neq \mathcal{I} \subseteq [Q]} (-1)^{|\mathcal{I}|+1} \mathbb{I}_\Omega^{(\mathcal{I}_\mathcal{R})} \big(p_\Omega^{(\mathcal{I}_\mathcal{C})}(\mathcal{H}^{(\mathcal{R})}) \nabla \delta V^{(\mathcal{I})} + \delta V^{(\mathcal{I})} \nabla p_\Omega^{(\mathcal{I}_\mathcal{C})}(\mathcal{H}^{(\mathcal{R})}) \big), \quad (29)$$

Proof. Fix q_{run} and $\boldsymbol{x}^{(\mathcal{C})}$ with $\boldsymbol{x}^{(1:Q)} \in \mathcal{X}_{\neq}^{Q}$.

Non-degeneracy: by Assumption 1 and $\sigma_\epsilon^2 > 0$, the joint covariance is positive definite, which Gaussian conditioning and marginalization preserve, so the laws of $(\boldsymbol{g}^{(q)})_{q \in \mathcal{I}_\mathcal{C}}$ given $\mathcal{H}^{(\mathcal{R})}$ ($\mathcal{I} \not\subseteq \mathcal{R}$) and of $\boldsymbol{g}^{(Q)}$ given $\mathcal{H}_g$ are non-degenerate Gaussians.

Frozen block: on the frozen draw, the generating vector $\boldsymbol{z}^{(\mathcal{R})} := (\boldsymbol{f}^{(1:Q)}, \boldsymbol{g}^{(\mathcal{R})})$ of $\mathcal{H}^{(\mathcal{R})}$ is the affine image $\boldsymbol{\mu}^{(\mathcal{R})} + L^{(\mathcal{R})}\tilde{\boldsymbol{\varepsilon}}$. In batch-position order, the rows of $\boldsymbol{\mu}^{(\mathcal{R})}$ and $L^{(\mathcal{R})}$ at fixed positions are built from kernel evaluations at $\boldsymbol{x}^{(\mathcal{R})}$ and the observed inputs alone, so they are constant in $\boldsymbol{x}^{(\mathcal{C})}$. In particular, every $\mathbb{I}_\Omega^{(\mathcal{I}_\mathcal{R})}$ is constant in $\boldsymbol{x}^{(\mathcal{C})}$.

C^1 **moments**: the rows at positions $\mathcal{C}$ and the conditional moments $(\boldsymbol{\mu}_{\mathcal{I}_\mathcal{C}}, \Sigma_{\mathcal{I}_\mathcal{C}})$ are C^1 in $\boldsymbol{x}^{(\mathcal{C})}$. This is because they are compositions of C^1 kernel evaluations including boundary points of $\mathcal{X}$ by Assumption 1 with inversion and Cholesky factorization, which are smooth on their blocks, a property preserved by the non-degeneracy step (Murray, 2016).

C^1 **orthant probability**: the Gaussian orthant probability $\Psi(\boldsymbol{\mu}, \Sigma) := \int_{(-\infty,0]^d} \mathcal{N}(\boldsymbol{g} \mid \boldsymbol{\mu}, \Sigma)\, d\boldsymbol{g}$, $d := |\mathcal{I}_\mathcal{C}|L$, is C^1 on $\{\Sigma \succ 0\}$ with the closed-form partials (see Proposition 1 of Marmin et al. (2015)). Thus, $p_\Omega^{(\mathcal{I}_\mathcal{C})}(\mathcal{H}^{(\mathcal{R})}) = \Psi(\boldsymbol{\mu}_{\mathcal{I}_\mathcal{C}}, \Sigma_{\mathcal{I}_\mathcal{C}})$ and $p_\Omega^{(Q)}(\mathcal{H}_g)$ are C^1 on the stated set by composition because a non-degenerate Gaussian charges every non-empty open set.

Factorization: $\mathbb{I}_\Omega^{(\mathcal{I})} = \mathbb{I}_\Omega^{(\mathcal{I}_\mathcal{R})} \mathbb{I}_\Omega^{(\mathcal{I}_\mathcal{C})}$ by the product form of $\mathbb{I}_\Omega^{(\mathcal{I})}$ with $\mathbb{I}_\Omega^{(\mathcal{I}_\mathcal{R})}$ being $\mathcal{H}^{(\mathcal{R})}$-measurable, so the pull-out property gives the factorization of (i). The case $\mathcal{I} \subseteq \mathcal{R}$ is included by the convention $\mathbb{I}_\Omega^{(\emptyset)} = 1$. Hence, with the three previous steps, it proves (i).

Local Lipschitzness: each $\delta V^{(\mathcal{I})}$ is an inclusion-exclusion term of Appendix C of Daulton et al. (2020), hence locally Lipschitz in $\boldsymbol{x}^{(\mathcal{C})}$ and Lebesgue-a.e. differentiable.

Conclusion: through the conditional-PoF factorization, the utility reads $u_{\mathrm{RB}}(\mathcal{H}^{(\mathcal{R})})$, a finite sum of products of a locally Lipschitz, a constant and a C^1 factor, so it is locally Lipschitz, differentiable outside the union of the $2^Q - 1$ Lebesgue-null sets. The product rule yields Eq. (29) with the constant $\mathbb{I}_\Omega^{(\mathcal{I}_\mathcal{R})}$ passing through untouched. A term with $\mathcal{I} \subseteq \mathcal{R}$ has $\nabla \delta V^{(\mathcal{I})} = \mathbf{0}$, since $\delta V^{(\mathcal{I})}$ reads the frozen $\boldsymbol{f}^{(\mathcal{R})}$ alone, and has the constant probability factor $p_\Omega^{(\emptyset)}(\mathcal{H}^{(\mathcal{R})}) = 1$, so it drops out of Eq. (29) and the sum is effectively over $\mathcal{I} \not\subseteq \mathcal{R}$. □

Remark 18 (Comparison with the Sigmoid Indicator). The baseline u_{sig} restores nonzero constraint derivatives, but for every fixed $\tau > 0$, it targets the wrong AF:

$$\alpha_{\mathrm{sig}} - \alpha_{\mathrm{hard}} = \mathbb{E}\Big[\Delta^{(Q)}\Big(\textstyle\prod_{l=1}^{L} \varsigma\big(-g_l^{(Q)}/\tau\big) - \mathbb{I}_\Omega^{(Q)}\Big)\Big] \neq 0 \text{ in general.} \quad (30)$$

$\tau \to 0^+$ removes the bias only by degenerating the constraint derivative. By contrast, $u_{\mathrm{RB}}(\mathcal{H})$ is exact at every $\boldsymbol{x}^{(1:Q)}$ and every lattice member with no smoothing and no limit (Theorem 9), and carries no tunable parameter.

A.6 Reduction to Univariate Gaussian CDFs under Sequential Greedy and Independence

The theorems above are quantified over the whole lattice and leave every $p_\Omega^{(\mathcal{I})}(\mathcal{H})$ a multivariate Gaussian CDF, which requires numerical integration. The conditions below cut that object down to what the implementation evaluates.

Assumption 3 (Sequential Greedy)**.** $\boldsymbol{x}^{(1:Q-1)}$ is pairwise distinct and fixed during AF optimization.

Assumption 4 (Objective-Constraint Independence)**.** The objective block $(f_m)_{m=1}^M$ and the constraint block $(g_l)_{l=1}^L$, together with their noises $\boldsymbol{\varepsilon}_y, \boldsymbol{\varepsilon}_c$, are independent *across the two blocks* a priori, so the joint posterior factorizes as $p(\boldsymbol{f}^{(1:Q)}, \boldsymbol{g}^{(1:Q)} \mid \mathcal{D}_N) = p(\boldsymbol{f}^{(1:Q)} \mid Y_N)\, p(\boldsymbol{g}^{(1:Q)} \mid \{\boldsymbol{c}_n\}_{n=1}^N)$.

Assumption 5 (Constraint Independence)**.** The L constraint processes $g_1, \ldots, g_L$, together with their noises $\varepsilon_{c,1}, \ldots, \varepsilon_{c,L}$, are mutually independent a priori.

All three hold throughout this section, of which Corollary 21 is the specialization of Theorems 9, 10, 12 and 17 to $\mathcal{H} = \mathcal{H}^{(1:Q-1)}$ under them, while the other statements below are new results. Our implementation adopts constraint-wise independent GPs. Write $\mathcal{X}_\circ \coloneqq \mathcal{X} \setminus \{\boldsymbol{x}^{(1)}, \ldots, \boldsymbol{x}^{(Q-1)}\}$ for the candidate inputs distinct from the running trials so that $\sigma_l^{(Q|1:Q-1)} > 0$ for every $l \in [L]$. Assumption 4 and Assumption 5 factorize the two orthant probabilities into products of univariate Gaussian CDFs,

$$p_\Omega^{(Q)}(\mathcal{H}_g) = \prod_{l=1}^{L} \Phi\Bigg(\underbrace{-\frac{\mu_l^{(Q|1:Q-1)}}{\sigma_l^{(Q|1:Q-1)}}}_{=t_l}\Bigg), \quad p_\Omega^{(Q)} = \prod_{l=1}^{L} \Phi\Big(-\frac{\mu_l^{(Q)}}{\sigma_l^{(Q)}}\Big), \tag{31}$$

where $\mu_l^{(Q|1:Q-1)}, (\sigma_l^{(Q|1:Q-1)})^2$ are computed by the conditional posterior given $\mathcal{H}^{(1:Q-1)}$, $\mu_l^{(Q)}$, $(\sigma_l^{(Q)})^2$ are the *unconditional* posterior, and t_l is the negated standardized violation.

Proposition 19 (Constraint Covariance Signs the Marginal-PoF Bias)**.** *For every $q \in [Q-1]$ and $l, l' \in [L]$, assume $\mathrm{Cov}\big[g_l^{(q)}, g_{l'}^{(Q)}\big] \geq 0$. Then $\alpha_{\mathrm{marg}} - \alpha_{\mathrm{hard}} = -\mathrm{Cov}[\Delta^{(Q)}, \mathbb{I}_\Omega^{(Q)}] \geq 0$ for every $\boldsymbol{x}^{(1:Q)} \in \mathcal{X}^Q$. If instead $\mathrm{Cov}\big[g_l^{(q)}, g_{l'}^{(Q)}\big] \leq 0$ for every such q, l, l', then $\alpha_{\mathrm{marg}} - \alpha_{\mathrm{hard}} \leq 0$.*

Proof. We sign the covariance for the stated case, the reversed one being verbatim. Fix $\boldsymbol{f}^{(1:Q)}$. $\mathbb{E}[\mathbb{I}_\Omega^{(Q)} \mid \mathcal{H}^{(\emptyset)}] = p_\Omega^{(Q)}$ is constant by Assumption 4, so $\mathrm{Cov}[\Delta^{(Q)}, \mathbb{I}_\Omega^{(Q)}] = \mathbb{E}[\mathrm{Cov}[\Delta^{(Q)}, \mathbb{I}_\Omega^{(Q)} \mid \mathcal{H}^{(\emptyset)}]]$, and the conditional law of $\boldsymbol{g}^{(1:Q)}$ given $\mathcal{H}^{(\emptyset)}$ is its posterior law. By Definition 6, increasing any coordinate $g_l^{(q)}$ for $q \in [Q-1]$ either leaves $Y_{N,\Omega}^{(1:Q-1)}$ unchanged or makes $\boldsymbol{g}^{(q)}$ infeasible, thereby removing $\boldsymbol{f}^{(q)}$ from $Y_{N,\Omega}^{(1:Q-1)}$. Because this weakly enlarges $\Delta^{(Q)}$, $\boldsymbol{g}^{(1:Q-1)} \mapsto \Delta^{(Q)}$ is coordinatewise non-decreasing, and $\boldsymbol{g}^{(Q)} \mapsto \mathbb{I}_\Omega^{(Q)}$ is coordinatewise non-increasing. Conditionally on $\mathcal{H}^{(\emptyset)}$, the claim is the Kahane inequality (Kahane, 1986) (see Proposition 5.1 of Biskup, 2017) applied against the independent coupling, and taking expectations completes the proof. □

Remark 20 (When the Sign Is Determined). The two cases of Proposition 19 require every entry of the cross-covariance block to carry the same sign. For mixed signs, no claim is made. The entries are posterior cross-covariances. The non-negative case is typical when the candidate is near the running trials relative to the data. Meanwhile, the reversed case requires all $Q-1$ cross-covariances to be negative simultaneously, so a single non-negative entry among them becomes increasingly likely as Q grows. Under Assumption 5, the block reduces to the $L(Q-1)$ covariances, so which case holds is a per-candidate runtime check rather than a property of the kernel class. Informally, whenever the non-negative case does hold, Proposition 19 says α_{marg} overestimates exactly at the candidates closest to the running trials, so it may be hard for α_{marg} to maintain the diversity within a batch.

Corollary 21 (Sequential-Greedy Univariate Form)**.** *For every $\boldsymbol{x}^{(1:Q)} \in \mathcal{X}^Q$:*

(i) ***Utility and AF Value****:*

$$u_{\mathrm{RB}} = B + \Delta^{(Q)}\, p_\Omega^{(Q)}(\mathcal{H}_g) \text{ a.s., so } \alpha_{\mathrm{hard}} = \mathbb{E}\big[B + \Delta^{(Q)}\, p_\Omega^{(Q)}(\mathcal{H}_g)\big]. \tag{32}$$

(ii) ***Variance Gap****:*

$$\mathbb{V}[u_{\mathrm{hard}}] - \mathbb{V}[u_{\mathrm{RB}}] = \mathbb{E}\big[\big(\Delta^{(Q)}\big)^2\, p_\Omega^{(Q)}(\mathcal{H}_g)\,(1 - p_\Omega^{(Q)}(\mathcal{H}_g))\big] \geq 0. \tag{33}$$

For S iid draws, $\mathbb{V}[\hat{\alpha}_{\mathrm{hard},S}] - \mathbb{V}[\hat{\alpha}_{\mathrm{RB},S}] = S^{-1}\,(\mathbb{V}[u_{\mathrm{hard}}] - \mathbb{V}[u_{\mathrm{RB}}]) \geq 0$.

(iii) ***Marginal-PoF Bias****: the marginal-PoF utility u_{marg} satisfies*

$$\alpha_{\mathrm{marg}} - \alpha_{\mathrm{hard}} = -\mathrm{Cov}\big[\Delta^{(Q)},\, p_\Omega^{(Q)}(\mathcal{H}_g)\big] = -\mathrm{Cov}\big[\Delta^{(Q)},\, \mathbb{I}_\Omega^{(Q)}\big]. \tag{34}$$

In particular, $\mathrm{MSE}[\hat{\alpha}_{\mathrm{RB},S}] < \mathrm{MSE}[\hat{\alpha}_{\mathrm{marg},S}]$ for every $S > \mathbb{V}[u_{\mathrm{RB}}]/(\alpha_{\mathrm{marg}} - \alpha_{\mathrm{hard}})^2$ if $\alpha_{\mathrm{marg}} \neq \alpha_{\mathrm{hard}}$.

(iv) ***Pathwise Candidate Gradient****: Fix fantasy samples. On $\mathcal{X}_\circ$, $p_\Omega^{(Q)}(\mathcal{H}_g)$ is C^1 in $\boldsymbol{x}^{(Q)}$ with values in $(0,1)$, and $\boldsymbol{x}^{(Q)} \mapsto u_{\mathrm{RB}}$ is locally Lipschitz, differentiable at Lebesgue-a.e. $\boldsymbol{x}^{(Q)}$, with*

$$\nabla_{\boldsymbol{x}^{(Q)}} u_{\mathrm{RB}} = p_\Omega^{(Q)}(\mathcal{H}_g)\, \nabla\Delta^{(Q)} + \Delta^{(Q)}\, \nabla p_\Omega^{(Q)}(\mathcal{H}_g),\ \nabla B = \mathbf{0}, \tag{35}$$

$$\nabla p_\Omega^{(Q)}(\mathcal{H}_g) = \sum_{l=1}^{L} \phi(t_l)\, \nabla_{\boldsymbol{x}^{(Q)}} t_l \prod_{l' \neq l} \Phi(t_{l'}), \tag{36}$$

a combination with strictly positive weights of L univariate terms.

Proof. (i) B and $\Delta^{(Q)}$ are $\mathcal{H}^{(1:Q-1)}$-measurable, so pulling them out of $\mathbb{E}[\cdot \mid \mathcal{H}^{(1:Q-1)}]$ in $u_{\mathrm{hard}} = B + \Delta^{(Q)}\mathbb{I}_\Omega^{(Q)}$ of Eq. (14) gives $u_{\mathrm{RB}}(\mathcal{H}^{(1:Q-1)}) = B + \Delta^{(Q)} p_\Omega^{(Q)}(\mathcal{H}^{(1:Q-1)})$ by Theorem 9 (i), and $p_\Omega^{(Q)}(\mathcal{H}^{(1:Q-1)}) = p_\Omega^{(Q)}(\mathcal{H}_g)$ by Assumption 4. The AF identity is Theorem 9 (ii) at $\mathcal{H} = \mathcal{H}^{(1:Q-1)}$.

(ii) $\mathbb{I}_\Omega^{(Q)} \mid \mathcal{H}^{(1:Q-1)}$ is Bernoulli $p_\Omega^{(Q)}(\mathcal{H}_g)$, so $\mathbb{V}[u_{\mathrm{hard}} \mid \mathcal{H}^{(1:Q-1)}] = (\Delta^{(Q)})^2 p_\Omega^{(Q)}(\mathcal{H}_g)(1 - p_\Omega^{(Q)}(\mathcal{H}_g))$ when $\mathcal{H} = \mathcal{H}^{(1:Q-1)}$ in Eq. (20), and Corollary 11 at $\mathcal{H} = \mathcal{H}^{(1:Q-1)}$ gives the estimator form.

(iii) This is Eq. (24) at $p_\Omega^{(Q)}(\mathcal{H}^{(1:Q-1)}) = p_\Omega^{(Q)}(\mathcal{H}_g)$. The latter follows from Corollary 16.

(iv) Under Assumption 3, Theorem 17 applies with $\mathcal{R} = [Q-1]$, so $\mathcal{H}^{(\mathcal{R})} = \mathcal{H}^{(1:Q-1)}$ and the stated open set is $\mathcal{X}_\circ$. With Eq. (31), $p \coloneqq p_\Omega^{(Q)}(\mathcal{H}_g)$ is C^1 in $\boldsymbol{x}^{(Q)}$ with values in $(0,1)$. On the frozen draw, B is built from the realized fantasies at the fixed $\boldsymbol{x}^{(1:Q-1)}$ alone (see Theorem 17), so $\nabla B = \mathbf{0}$. Meanwhile, $\Delta^{(Q)}$ is locally Lipschitz by the step at $\mathcal{I} = \{Q\}$. Hence, $u_{\mathrm{RB}}(\mathcal{H}^{(1:Q-1)}) = B + \Delta^{(Q)} p$ is locally Lipschitz, differentiable at Lebesgue-a.e. $\boldsymbol{x}^{(Q)}$, and the product rule gives Eq. (35). The chain rule on Eq. (31) gives Eq. (36), with $\phi(t_l)\prod_{l' \neq l} \Phi(t_{l'}) > 0$ as $\phi > 0$ and $\Phi \in (0,1)$.

□

The gradient behavior can also be analyzed.

Proposition 22 (Exact Log-PoF Gradient)**.** *On $\mathcal{X}_\circ$, the exact log-gradient of the candidate PoF is*

$$\nabla_{\boldsymbol{x}^{(Q)}} \log p_\Omega^{(Q)}(\mathcal{H}_g) = \sum_{l=1}^{L} \frac{\phi(t_l)}{\Phi(t_l)} \nabla_{\boldsymbol{x}^{(Q)}} t_l. \tag{37}$$

Proof. Write $\lambda(t) \coloneqq \phi(t)/\Phi(t)$, so Eq. (31) reads $p_\Omega^{(Q)}(\mathcal{H}_g) = \prod_{l=1}^{L} \Phi(t_l)$. By Corollary 21 (iv), $p_\Omega^{(Q)}(\mathcal{H}_g) \in (0,1)$ and each t_l is C^1 on $\mathcal{X}_\circ$, so $\log p_\Omega^{(Q)}(\mathcal{H}_g) = \sum_{l=1}^{L} \log \Phi(t_l)$ is C^1 there, and the chain rule with $\Phi' = \phi$ gives $\nabla \log p_\Omega^{(Q)}(\mathcal{H}_g) = \sum_{l=1}^{L} \lambda(t_l)\nabla t_l$, i.e., Eq. (37). □

Lemma 23 (Tail Behavior of the Exact Log-PoF Score)**.** *The coefficient $\phi(t)/\Phi(t)$ of Eq. (37), as a function of $t \in \mathbb{R}$, satisfies $\phi(t)/\Phi(t) = -t + O(|t|^{-1})$ as $t \to -\infty$ and $\phi(t)/\Phi(t) \to 0$ as $t \to +\infty$.*

Proof. Set $t \to -\infty$. The inverse Mills ratio gives $\lambda(t) = \phi(-t)/(1 - \Phi(-t)) = -t + O(|t|^{-1})$. The second claim follows since $\Phi(t) \to 1$ and $\phi(t) \to 0$ as $t \to +\infty$. □

Remark 24 (Comparison of the Log-PoF Score and the Sigmoid Relaxation). The log-gradients of the sigmoid approach are $\partial_t \log \varsigma(t) = 1 - \varsigma(t)$ at $t \coloneqq -g_l^{(Q)}/\tau$, whereas our approach has $\lambda(t_l) \coloneqq \phi(t_l)/\Phi(t_l)$ at t_l. The sigmoid's log-gradients vanish as $t \to \infty$ and saturate at 1 as $t \to -\infty$. Thus, on the infeasible tail, its gradients approach the violation-independent $\nabla g_l^{(Q)}/\tau$ and no longer separate a mild violation from a severe one. The coefficient of Lemma 23 instead grows unboundedly as $\lambda(t_l) = -t_l + O(|t_l|^{-1})$ with no temperature parameter, while still vanishing on the feasible tail.

Remark 25 (Role of the Three Assumptions). The three conditions jointly make the univariate form the operative estimator. Each is sufficient, and dropping one leaves nothing to enforce its conclusion rather than negating it.

1. Without Assumption 3, $\boldsymbol{x}^{(1:Q-1)}$ is no longer globally fixed during joint batch optimization. However, whenever a subset $\mathcal{R}$ of batch points is held fixed, its contribution can still be cached within the corresponding block subproblem. Although this requires evaluating the $O(2^{Q-|\mathcal{R}|})$ inclusion-exclusion terms, the assumption does not affect the validity of our RB results,
2. Without Assumption 4, the conditioning on $\boldsymbol{f}^{(1:Q)}$ inside $\mathcal{H}^{(1:Q-1)}$ shifts the constraint mean and couples the L constraints even when they are a priori independent. Since each of them explains part of the same objective draw, $p_\Omega^{(Q)}(\mathcal{H}^{(1:Q-1)})$ need not factorize and is in general an L-dimensional Gaussian CDF, and
3. Without Assumption 5, the L constraints are correlated a priori, and $p_\Omega^{(Q)}(\mathcal{H}_g)$ is again in general an L-dimensional Gaussian CDF.

A correlated objective-constraint block needs a matrix-valued multi-output kernel with a non-zero cross-block, of which the latent Kronecker structure of Lin et al. (2025); Watanabe (2025) is a computationally tractable choice.

B Detailed Experimental Setup

Baseline Methods. Optuna `GPSampler` uses Matérn-5/2 kernel with gradient-based MAP estimation for the kernel hyperparameters to achieve good performance. The AF uses the number of fantasy samples $S = 128$ following (Daulton et al., 2020). Optuna optimizes the AF using L-BFGS-B with multi-start local searches (Irie et al., 2025) where Optuna restarts L-BFGS-B 10 times. Since Optuna `GPSampler` currently supports the sigmoid approach, we directly use it in our experimental pipeline, and all our implementation, except for the PoF term computation, shares the `GPSampler` pipeline for fair comparison.

Among many TPE Implementation variants including Hyperopt (Bergstra et al., 2015) and HpBandSter (Falkner et al., 2018), Optuna mostly follows the recommended setup detailed by Watanabe (2023). The multi-objective optimization part of TPE follows MOTPE (Ozaki et al., 2020; 2022) and the constrained part follows joint c-TPE (Watanabe & Irie, 2026), which is the joint modeling variant of c-TPE (Watanabe & Hutter, 2023).

Implementation details. All experiments are implemented in Python 3.12.12 using Optuna v5.0.0, OptunaHub v0.5.0, PyTorch 2.13.0+cpu, and SciPy 1.18.1. We use the OptunaHub registry (Ozaki et al., 2026) at commit `abac65b96fef633bba86b85544da12ee473ffdd4`.

Machine specifications. All experiments are run on a single machine with an Intel® Core™ Ultra 7 155H CPU (22 logical CPUs) and 16 GB RAM, running Arch Linux.