LEARN FEASIBILITY ONCE, OPTIMIZE ALL OBJECTIVES: DERIVATIVE-FREE DIFFUSION MODELS FOR CHANCE-CONSTRAINED PROGRAMMING

Anonymous authors
Paper under double-blind review

ABSTRACT

Chance-constrained programs (CCPs) optimize decisions under uncertainty by limiting the probability of constraint violation. Despite advances in traditional and learning-based approaches, optimizing non-convex or non-smooth objectives and adapting to different objectives under fixed chance constraints remain challenging. In this paper, we propose a **D**erivative-free **D**iffusion-based framework that **D**isentangles constraint modeling from objective optimization, termed **D**^{3}**Opt**. We learn the chance-feasible structure once, independently of any particular objective, by training a risk-conditioned diffusion model solely on constraint-filtered decisions and freezing it as a reusable prior for post-specified objectives. At inference time, we propose an annealed, particle-based Feynman–Kac correction along the frozen reverse diffusion process to optimize post-specified objectives using only function evaluations. This enables derivative-free optimization of non-convex and non-smooth objectives without objective-specific retraining. We prove that the correction preserves feasibility when this property holds for the frozen prior, and derive an optimization-error bound separating learned-prior coverage, finite-particle approximation, and finite-temperature effects. Experiments on linear Gaussian CCPs, objective-transfer tasks, and chance-constrained economic dispatch demonstrate effective optimization across smooth and non-smooth objectives, including non-convex cases, and objective generalization under fixed chance constraints without retraining.

1 INTRODUCTION

Chance-constrained programming provides a principled framework for decision making under uncertainty by requiring constraints to hold with a prescribed probability. It has applications in economic dispatch (Wang et al., 2017; Yang et al., 2021), optimal power flow (Bienstock et al., 2014), vehicle routing (Dinh et al., 2018), and production planning (Zhang et al., 2023). Despite extensive research, evaluating probabilistic constraints and optimizing over the resulting non-convex feasible regions can remain computationally challenging (Nemirovski & Shapiro, 2007; Pagnoncelli et al., 2009). In many applications, the uncertainty model and reliability requirements remain fixed while operational objectives or preferences change (Cao et al., 2013; Ren & Wang, 2026). These tasks share the same chance-feasible structure, even though they favor different decisions. This observation raises the first question:

Q1: Can feasibility be learned once and reused across objectives?

Learning-based approaches use data to guide optimization, and diffusion-based methods learn decision distributions that support constrained solution generation (Kong et al., 2025; Zhang et al., 2025; Ding et al., 2025). However, the extent to which these distributions can be reused across objectives depends on their construction. GGDOpt (Zhang et al., 2025), for example, trains on solutions of objective-dependent restricted problems, so its training data are tailored to a particular objective rather than constructed solely from constraint information. To remove this dependence, we separate learning *where feasible decisions lie* from deciding *which feasible decisions are preferred*. We generate training decisions using only constraint information and train a risk-conditioned diffu-

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.

sion model without evaluating the objective. Once trained, the model is frozen and reused across post-specified objectives under fixed chance constraints, without objective-specific data generation.

Reusing the learned constraint prior across objectives also requires handling the different optimization challenges these objectives introduce. In practical applications, objectives may be non-convex, non-smooth, or accessible only through function evaluations (Masutomi et al., 2013; Guesmi et al., 2020). Since guidance based on objective gradients or Hessians cannot cover all such cases, a second question follows:

Q2: How can the frozen constraint prior support optimization of non-convex and non-smooth objectives using only function evaluations?

We introduce an annealing path that progressively strengthens the objective preference along the frozen reverse diffusion process and implement it through particle-based Feynman–Kac correction (FKC) (Del Moral, 2004; Skreta et al., 2025; Singhal et al., 2025). The correction reweights and resamples particles using objective values evaluated at denoised decision predictions, concentrating the search toward low-objective regions without requiring objective gradients or Hessians. The same frozen model can therefore support derivative-free optimization across different objectives.

By combining the reusable constraint prior with the inference-time Feynman–Kac correction, we obtain **D^{3}Opt**, a **D**erivative-free **D**iffusion-based framework that **D**isentangles constraint modeling from objective optimization. Our main contributions are:

- We propose D^{3}Opt, which learns a reusable risk-conditioned diffusion prior solely from constraint information and supports objective generalization without retraining.
- We develop an annealed, particle-based Feynman–Kac correction that optimizes post-specified non-convex and non-smooth objectives using only function evaluations.
- We establish feasibility preservation under an almost-surely feasible prior and derive an error bound separating prior coverage, finite-particle, and finite-temperature effects.
- We demonstrate solution quality and objective generalization on linear Gaussian CCPs, objective-transfer tasks, and chance-constrained economic dispatch.

2 Related Work

Chance-constrained optimization. Classical methods for CCPs include deterministic reformulations, conservative convex approximations, and sample average approximation (SAA) (Nemirovski & Shapiro, 2007; Pagnoncelli et al., 2009). The non-convexity and discontinuity of empirical chance constraints motivate mixed-integer, difference-of-convex, and zero–one optimization methods (Wang et al., 2023; Zhou et al., 2024). Derivative-free solvers such as differential evolution and mesh adaptive direct search also accommodate non-smooth objectives (Storn & Price, 1997; Audet & Dennis, 2006). These methods optimize a specified objective for each problem instance rather than learning a prior that can be reused across objectives.

Diffusion models for constrained optimization. Diffusion-based methods learn decision distributions and use generated samples to search under constraints (Ding et al., 2025; Chen et al., 2024). DiffOPT (Kong et al., 2025) represents unknown constraints through a learned data distribution and combines it with an objective-induced Boltzmann factor. It uses gradient-guided sampling for differentiable objectives and iterative importance sampling for non-differentiable ones. GGDOpt (Zhang et al., 2025) addresses CCPs by training on solutions of objective-dependent restricted problems and applying first- or second-order guidance at inference. Our framework learns a risk-conditioned prior from constraint-only data and applies annealed Feynman–Kac particle correction to the frozen sampler, adapting it to new objectives using objective values without retraining.

3 Preliminaries

3.1 Problem Definition

We consider the chance-constrained programming:

$$\min_{x\in\mathcal{X}} f(x) \qquad \text{s.t.} \qquad \mathbb{P}_{\xi\sim\mathbb{P}}\big(g(x,\xi)\le 0\big)\ge 1-\rho, \tag{1}$$

where $\mathcal{X} \subseteq \mathbb{R}^d$ is the decision domain, $f : \mathcal{X} \to \mathbb{R}$ is the objective, ξ represents uncertainty, and $\rho \in (0, 1)$ is the admissible violation probability. The chance-feasible set is

$$\mathcal{X}_\rho = \left\{ x \in \mathcal{X} : \mathbb{P}_{\xi \sim \mathbb{P}}\big(g(x, \xi) \leq 0\big) \geq 1 - \rho \right\}. \tag{2}$$

Following the optimization-as-sampling perspective (Kong et al., 2025; Zhang et al., 2025), let $p_{\mathrm{F}}(\cdot \mid \rho)$ be a reference distribution supported on $\mathcal{X}_\rho$. For inverse temperature $\beta \geq 0$, we define

$$\pi_\beta(x \mid \rho) = \frac{p_{\mathrm{F}}(x \mid \rho) e^{-\beta f(x)}}{Z_\beta(\rho)}, \qquad Z_\beta(\rho) = \int_{\mathcal{X}} p_{\mathrm{F}}(x \mid \rho) e^{-\beta f(x)}\, dx, \tag{3}$$

assuming $0 < Z_\beta(\rho) < \infty$. Under suitable regularity conditions and positive reference mass in neighborhoods of feasible global minimizers, π_β concentrates on these minimizers as $\beta \to \infty$.

3.2 Diffusion Model

Score-based diffusion models learn to reverse a stochastic noising process using the score of its marginals (Song et al., 2021b). We write the learned reverse process in generative time as:

$$dX_t = b_{\theta,t}(X_t, \rho)\, dt + \sigma_t\, dB_t, \qquad X_0 \sim p_{\mathrm{ref}}, \tag{4}$$

where $b_{\theta,t}$ contains the learned score contribution. We denote its path law by $\mathbb{Q}_\theta(\cdot \mid \rho)$ and its terminal distribution by $p_\theta(\cdot \mid \rho)$. Beyond sampling from the pretrained distribution itself, a growing line of work studies inference-time tilting toward an unnormalized target

$$\pi_{\theta,\beta}(x \mid \rho) \propto p_\theta(x \mid \rho) \exp[\beta r(x)].$$

Recent methods construct such information through value functions, denoised predictions, or lookahead rewards (Li et al., 2025; Kim et al., 2026). Other approaches instead control the proposal dynamics or combine transport with particle reweighting (Ren et al., 2026; Skreta et al., 2025).

For optimization, however, a fixed tilt is generally not the most natural viewpoint. As discussed in Section 3.1, approaching the global optimum corresponds to the low-temperature limit $\beta \to \infty$, which motivates an annealed sequence $0 = \beta_0 < \beta_1 < \cdots < \beta_K$ of progressively sharper target distributions. This creates a sequential change-of-measure problem: the sampler must gradually redistribute probability mass from the pretrained diffusion prior toward increasingly objective-favored regions. FKC provides a method for this purpose by representing the sequence of target changes.

3.3 Feynman–Kac Correction

Feynman–Kac change of measure. Let $\mathbb{Q}$ denote the path law of a baseline diffusion. Feynman–Kac (FK) transforms modify this law by assigning nonnegative multiplicative weights to trajectories (Del Moral, 2004; Lelièvre et al., 2010). Given a terminal potential G and a running potential V_t, the transformed path law takes the form

$$\frac{d\mathbb{Q}^{G,V}}{d\mathbb{Q}}(X_{0:T}) \propto \exp\left(- \int_0^T V_t(X_t)\, dt \right) G(X_T).$$

Thus, Feynman–Kac correction is fundamentally a change of measure on the baseline dynamics. It may be realized through a drift modification when the associated backward information function is available, or through particle reweighting without modifying the baseline dynamics. Further continuous-time details are provided in Appendix A.

Feynman–Kac correction with particles. Instead of modifying the diffusion drift, FKC can keep the baseline dynamics unchanged and adjust particle probabilities through multiplicative weights. This yields a transport–diffusion–reweighting view of the FK evolution. Sequential Monte Carlo (SMC) approximates this process by alternating particle propagation, reweighting, and resampling (Del Moral, 2004). This particle formulation underlies recent FK-based inference-time steering methods for diffusion models (Skreta et al., 2025; Ren et al., 2026; Ou et al., 2026).

For our method, consider a discretization $0 = t_0 < \cdots < t_K = T$ and denote the baseline Markov kernels by $M_k := P_{t_{k-1}, t_k}$. Given nonnegative incremental FK potentials G_k :

$\mathcal{X} \times \mathcal{X} \to \mathbb{R}_+$, the FK correction reweights each baseline trajectory $X_{0:K}$ by the multiplicative factor $\prod_{k=1}^{K} G_k(X_{k-1}, X_k)$. Thus, the pretrained diffusion kernels retain their role as the transport mechanism, whereas the FK potentials redistribute probability mass toward a prescribed sequence of tilted distributions. The formal discrete-time FK measures and their connection to the continuous formulation are given in Appendix A.

4 D^{3}OPT

We introduce D^{3}Opt, a derivative-free, diffusion-based framework that disentangles constraint modeling from objective optimization. Section 4.1 describes how D^{3}Opt learns a reusable risk-conditioned diffusion prior from constraint-only data. Section 4.2 presents the annealed, particle-based Feynman–Kac correction for optimizing post-specified objectives using only function evaluations. Section 4.3 establishes feasibility preservation under the frozen prior and analyzes the resulting optimization error.

4.1 OBJECTIVE-FREE FEASIBLE DIFFUSION MODEL

Constraint-only data generation. To decouple constraint learning from objective optimization, we construct the training dataset without evaluating any objective function or solving objective-dependent optimization problems. Candidate decisions are generated and retained using only the decision domain and chance-constraint information. The resulting dataset can therefore be reused when the objective changes while the constraints remain fixed.

For each problem, we choose an objective-independent proposal distribution ν on the deterministic decision domain $\mathcal{X}$. We independently draw a candidate $X \sim \nu$ and a target risk level $\rho \sim \omega$, where ω is a density on $\mathcal{R} = [\rho_{\min}, \rho_{\max}]$. The proposal distribution may differ across problems to reflect their decision domains, but remains fixed across objectives for the same problem.

We assess each candidate using two independent scenario sets, Ξ_{lab} and Ξ_{val}. These provide empirical estimates of the constraint-satisfaction probability $c(x) = \mathbb{P}_{\xi\sim\mathbb{P}}(g(x,\xi) \le 0)$:

$$\widehat{c}_{\mathrm{lab}}(x) = \frac{1}{N_{\mathrm{lab}}} \sum_{s=1}^{N_{\mathrm{lab}}} \mathbb{I}\{g(x, \xi_{\mathrm{lab}}^{(s)}) \le 0\}, \qquad \widehat{c}_{\mathrm{val}}(x) = \frac{1}{N_{\mathrm{val}}} \sum_{s=1}^{N_{\mathrm{val}}} \mathbb{I}\{g(x, \xi_{\mathrm{val}}^{(s)}) \le 0\}.$$

For a chance constraint, the indicator requires all components of g to be non-positive. We retain (X, ρ) when

$$A_N(X, \rho) = \mathbb{I}\left\{\widehat{c}_{\mathrm{lab}}(X) \ge 1 - \rho,\ \widehat{c}_{\mathrm{val}}(X) \ge 1 - \rho\right\} = 1.$$

Repeating this procedure produces $\mathcal{D}_F = \{(x^{(i)}, \rho^{(i)})\}_{i=1}^{N_D}$, where each decision is paired with its sampled target risk. Conditional on the scenario sets, accepted decisions at a fixed risk level follow

$$p_{F,N}(x \mid \rho) \propto \nu(x) A_N(x, \rho).$$

Both ν and A_N are independent of the objective. The diffusion model thus learns from constraint-filtered decisions without a training preference for any particular objective.

Conditional diffusion training. We instantiate the score-based diffusion framework introduced in Section 3.2 using a variance-preserving DDPM with the standard noise-prediction parameterization. For each objective-free feasible sample $(X_0, \rho) \in \mathcal{D}_F$, the forward process is

$$X_t = \sqrt{\bar{\alpha}_t} X_0 + \sqrt{1 - \bar{\alpha}_t}\, \epsilon, \qquad \epsilon \sim \mathcal{N}(0, I).$$

We train a risk-conditioned noise predictor $\epsilon_\theta(X_t, t, \rho)$ using the standard denoising objective

$$\mathcal{L}_{\mathrm{diff}}(\theta) = \mathbb{E}_{(X_0,\rho),t,\epsilon}\left[\|\epsilon - \epsilon_\theta(X_t, t, \rho)\|_2^2\right].$$

The resulting numerical reverse process defines conditional Markov kernels $M_k^\theta(x_{k-1}, dx_k \mid \rho)$, $k = 1, \ldots, K$, with intermediate marginals $q_k^\theta(\cdot \mid \rho)$ and terminal distribution $p_\theta(\cdot \mid \rho)$ approximating the objective-free feasible distribution learned from $\mathcal{D}_F$.

The objective function f is absent from data generation, diffusion training, checkpoint selection, and model parameters. After training, the conditional reverse process is frozen and used as the baseline transport for the inference-time Feynman–Kac correction introduced in the next subsection.

4.2 Derivative-Free Feynman–Kac Correction

After training, the diffusion model is frozen. We then introduce a post-specified objective $f : \mathcal{X} \to \mathbb{R}$, which may be non-convex, non-smooth, or available only through function evaluations.

Annealed derivative-free FKC. Consider the generative-time discretizations $0 = t_0 < t_1 < \cdots < t_K = T$. For an intermediate particle state X_k, let $\widehat{X}_k := \mathsf{D}_{\theta,k}(X_k, \rho)$ denote the terminal decision predicted by the frozen conditional diffusion model, with $\widehat{X}_K = X_K$ at the final step. Since the objective is defined on the terminal decision, we evaluate the current particle through $f(\widehat{X}_k)$. To progressively concentrate the particle system toward low-objective regions, we introduce an inverse-temperature schedule $0 = \beta_0 \le \beta_1 \le \cdots \le \beta_K = \beta_{\max}$ and define the incremental FK potential

$$G_k(X_{k-1}, X_k; \rho) \triangleq \exp\Big[\beta_{k-1} f(\widehat{X}_{k-1}) - \beta_k f(\widehat{X}_k)\Big], \qquad k = 1, \dots, K. \tag{5}$$

This potential requires only evaluations of the objective, in particular, no gradient or Hessian of f is needed. Moreover, the incremental potentials telescope along the reverse diffusion trajectory

$$\prod_{k=1}^{K} G_k(X_{k-1}, X_k; \rho) = \exp\Big[\beta_0 f(\widehat{X}_0) - \beta_K f(\widehat{X}_K)\Big] = \exp\left[-\beta_K f(X_K)\right],$$

where $\beta_0 = 0$ and $\widehat{X}_K = X_K$. Hence the terminal FK marginal is $\pi_{\theta,\beta_K}(x \mid \rho) \propto p_\theta(x \mid \rho) \exp[-\beta_K f(x)]$. Thus, the intermediate denoised predictions determine how particle mass is redistributed during sampling, but do not alter the desired terminal change of measure. The chance constraint does not enter the FK potential through an explicit penalty, and the risk level ρ instead enters through the frozen conditional diffusion model.

Interacting-particle approximation. We approximate the FK flow using M interacting particles. Starting from $X_0^{(m)} \sim p_{\mathrm{ref}}$ and $w_0^{(m)} = 1/M$, each particle is propagated through the frozen conditional reverse diffusion $X_k^{(m)} \sim M_k^\theta(X_{k-1}^{(m)}, \cdot \mid \rho)$, $k = 1, \dots, K$. We use a prescribed inverse-temperature schedule to progressively increase the influence of the post-specified objective:

$$\beta_k = \beta_{\max}\,(k/K)^\nu, \qquad k = 0, \dots, K, \quad \nu > 0. \tag{6}$$

Here, $\beta_{\max}$ controls the terminal objective tilt, while ν determines how rapidly the objective influence is introduced along the generative trajectory. The schedule is fixed before sampling and satisfies $\beta_0 = 0$ and $\beta_K = \beta_{\max}$. Using the incremental potential in equation 5, the particle weights are updated as

$$\widetilde{w}_k^{(m)} = w_{k-1}^{(m)} G_k(X_{k-1}^{(m)}, X_k^{(m)}; \rho), \qquad w_k^{(m)} = \widetilde{w}_k^{(m)} \Big/ \sum_{j=1}^{M} \widetilde{w}_k^{(j)}.$$

Although the temperature schedule is prescribed, resampling is performed adaptively according to the resulting weight degeneracy. Since the particle weights are normalized, their effective sample size is $\mathrm{ESS}_k = 1/\sum_{m=1}^{M} (w_k^{(m)})^2$. Whenever $\mathrm{ESS}_k < \tau_{\mathrm{res}} M$, with $\tau_{\mathrm{res}} \in (0, 1]$, we apply systematic resampling according to $\{w_k^{(m)}\}_{m=1}^M$ and reset all weights to $1/M$. Otherwise, the normalized weights are retained for the next propagation step. Thus, the prescribed temperature schedule controls the progressive objective tilt, while ESS is used only to determine when particle resampling is necessary.

4.3 Feasibility Preservation and Error Analysis

We next characterize the theoretical consequences of our method. We first show that, at the ideal measure level, FKC preserves any almost-sure feasibility already encoded by the frozen diffusion prior. We then analyze how finite-temperature concentration, approximation of the ideal feasible distribution, and finite-particle simulation affect the resulting optimizer.

Feasibility preservation under FKC. A key consequence of the proposed objective–constraint disentanglement is that the objective-only FKC cannot create probability mass outside the support of the frozen diffusion prior. Consequently, if the frozen diffusion is chance-feasible almost surely, the FK-tilted distribution remains chance-feasible almost surely, FKC itself does not introduce infeasibility. Let $\mathbb{Q}_\theta(\cdot \mid \rho)$ denote the path law of the frozen conditional reverse diffusion, and let $\mathbb{Q}_\theta^{\mathrm{FK}}(\cdot \mid \rho)$ denote the FK transform induced by the incremental potentials in equation 5.

Proposition 1 (Feasibility preservation under FKC)**.** *Fix $\rho \in (0,1)$ and suppose that the Feynman–Kac path weight is nonnegative, integrable, and has a strictly positive normalizing constant, then $\mathbb{Q}_\theta^{\mathrm{FK}}$ is absolutely continuous with respect to $\mathbb{Q}_\theta(\cdot \mid \rho)$, i.e. $\mathbb{Q}_\theta^{\mathrm{FK}}(\cdot \mid \rho) \ll \mathbb{Q}_\theta(\cdot \mid \rho)$. Consequently, the corresponding terminal marginal satisfies*

$$\pi_{\theta,\beta_K}(\cdot \mid \rho) \ll p_\theta(\cdot \mid \rho). \tag{7}$$

In particular, if the frozen diffusion is chance-feasible almost surely, i.e. $p_\theta(\mathcal{X}_\rho \mid \rho) = 1$, then the FK-corrected terminal distribution also satisfies $\pi_{\theta,\beta_K}(\mathcal{X}_\rho \mid \rho) = 1$.

Proof of Proposition 1 is given in Appendix B.1. For the objective-only tilt, $\frac{d\pi_{\theta,\beta_K}(\cdot|\rho)}{dp_\theta(\cdot|\rho)}(x) = \frac{\exp[-\beta_K f(x)]}{Z_{\theta,\beta_K}(\rho)}$, whenever $0 < Z_{\theta,\beta_K}(\rho) < \infty$. Hence FKC only redistributes mass within the support of the frozen prior and cannot create mass on baseline null sets. The condition $p_\theta(\mathcal{X}_\rho \mid \rho) = 1$ is an ideal support assumption and need not hold exactly under the relaxed finite-scenario construction and finite-data diffusion approximation in Section 4.1. Proposition 1 isolates the effect of FK correction by showing that FKC does not introduce support outside that of the frozen prior.

Approximation and optimization error. Proposition 1 shows that the objective-only Feynman–Kac correction preserves feasibility already encoded by the frozen diffusion prior. We next quantify the practical optimization error arising from incomplete coverage of the learned feasible distribution, finite inverse temperature, and finite-particle approximation.

For a fixed risk level ρ, define $f_\rho^\star \triangleq \inf_{x \in \mathcal{X}_\rho} f(x)$, and $f_{\theta,\rho}^\star \triangleq \operatorname{ess\,inf}_{x \sim p_\theta(\cdot|\rho)} f(x)$ is the essential infimum of f taken with respect to $p_\theta(\cdot \mid \rho)$. We further define the learned-prior coverage gap

$$\Delta_{\theta,\rho} \triangleq f_{\theta,\rho}^\star - f_\rho^\star.$$

Under almost-sure feasibility of the frozen diffusion, $\Delta_{\theta,\rho} \geq 0$. In particular, $\Delta_{\theta,\rho} = 0$ whenever the learned prior assigns positive probability to every arbitrarily near-optimal feasible objective sublevel. For any $\epsilon > 0$, define the near-optimal mass of the learned prior by

$$a_{\theta,\epsilon} \triangleq p_\theta\big(f(X) \leq f_{\theta,\rho}^\star + \epsilon \,\big|\, \rho\big)\,. \tag{8}$$

By the definition of the essential infimum, $a_{\theta,\epsilon} > 0$ for every $\epsilon > 0$. Let η_K denote the exact normalized terminal FK marginal and $\eta_K^M = \sum_{m=1}^M w_K^{(m)} \delta_{X_K^{(m)}}$ its normalized particle approximation.

Assumption 1 (Regularity for practical error analysis)**.** *Fix $\rho \in (0,1)$. Assume $p_\theta(\mathcal{X}_\rho \mid \rho) = 1$, $f_{\theta,\rho}^\star > -\infty$. For every fixed finite inverse-temperature schedule $0 = \beta_0 \leq \cdots \leq \beta_K < \infty$, assume that the associated FK model is well defined and that its M-particle approximation satisfies $(\mathbb{E}[|\eta_K^M(f) - \eta_K(f)|^2])^{1/2} \leq C_{K,f}/\sqrt{M}$ for some finite $C_{K,f}$ independent of M.*

The first condition is the feasibility premise of Proposition 1. The particle-error condition assumes the standard fixed-horizon $O(M^{-1/2})$ SMC rate, which holds for bounded observables and extends to unbounded test functions under suitable moment and stability conditions (Del Moral, 2004; Hu et al., 2011; Whiteley, 2012). Importantly, no differentiability, smoothness, or Hessian information of the post-specified objective f is required.

Proposition 2 (Practical optimization error)**.** *Under Assumption 1, for every finite $\beta_K > 0$, there exists a finite constant $C_{K,f}$, independent of M, such that*

$$\big|\mathbb{E}[\eta_K^M(f)] - f_\rho^\star\big| \leq \underbrace{\Delta_{\theta,\rho}}_{\textit{coverage error}} + \underbrace{\frac{C_{K,f}}{\sqrt{M}}}_{\textit{finite particles}} + \underbrace{\inf_{\epsilon>0}\left\{\epsilon + \frac{1}{\beta_K}\log\frac{1}{a_{\theta,\epsilon}}\right\}}_{\textit{finite inverse temperature}}. \tag{9}$$

Table 1: Comparison on the linear Gaussian chance-constrained problem equation 11 with $n = 8$ and $\rho = 0.1$. Lower objective values are better.

Method	Repeat	FvalMean ↓	FvalStd ↓	FvalMedian ↓
SOC-CVX (Grant et al., 2008) **(Global Reference)**	1	-0.6586	0	-0.6586
SAA-CVaR (Nemirovski & Shapiro, 2007)	100	-0.5893	0.0248	-0.5869
SAA-MIP (Pagnoncelli et al., 2009)	100	-0.6281	0.0157	-0.6318
SAA-PDCA (Wang et al., 2023)	100	-0.6389	0.0314	-0.6408
SAA-SNSCO (Zhou et al., 2024)	100	0.8051	3.4014	-0.6371
GGDOpt First-order (Zhang et al., 2025)	100	-0.6483	0.0051	-0.6488
GGDOpt Second-order (Zhang et al., 2025)	100	-0.6491	0.0056	-0.6503
DiffOPT (Kong et al., 2025)	100	-0.5614	0.0252	-0.5608
D^{3}Opt w/o FKC	100	-0.1172	0.0711	-0.1247
D^{3}Opt	100	$\mathbf{-0.6519}$	$\mathbf{0.0015}$	$\mathbf{-0.6517}$

Proof can be found in Appendix B.2. In practice, the SMC procedure returns a weighted terminal particle population, from which we output the one with the smallest objective value. Since $f(\widehat{X}_K^M) := \min_{1\le m\le M} f(X_K^{(m)}) \le \eta_K^M(f)$, the bound on $\mathbb{E}[\eta_K^M(f)]$ also provides an upper bound on the expected objective gap of the returned best particle. Moreover,

$$\limsup_{\beta_K\to\infty} \limsup_{M\to\infty} \left|\mathbb{E}[\eta_K^M(f)] - f_\rho^\star\right| \le \Delta_{\theta,\rho}. \tag{10}$$

In particular, when $\Delta_{\theta,\rho} = 0$, the infinite-particle limit followed by the low-temperature limit recovers the global chance-constrained optimal value in expectation. The analysis experiments are presented in Appendix D.2.

5 Experiments

We evaluate D^{3}Opt on linear Gaussian optimization (Section 5.1), objective transfer to smooth and non-smooth objectives (Section 5.2), and economic dispatch with valve-point effects (Section 5.3). Unless stated otherwise, stochastic methods are evaluated over 100 runs. D^{3}Opt uses a U-Net score model (Ronneberger et al., 2015) and stochastic DDIM sampling (Song et al., 2021a), with the DDIM stochasticity parameter fixed at $\eta = 1$ throughout the reported experiments. Implementation details and robust waveform experiments are provided in Appendices C and D.3. All experiments use one NVIDIA RTX 5090 GPU.

5.1 Linear Gaussian Chance-Constrained Optimization

We first consider the standard linear Gaussian chance-constrained optimization problem:

$$\min_{\mathbf{x}\in\mathbb{R}^n} \quad \frac{1}{2}\mathbf{x}^\top\mathbf{x} + \mathbf{b}^\top\mathbf{x}, \qquad \text{s.t.} \quad \mathbb{P}_{\mathbf{c}}\left\{\mathbf{c}^\top\mathbf{x} + d \ge 0\right\} \ge 1-\rho, \tag{11}$$

where $\mathbf{c} \sim \mathcal{N}(\bar{\mathbf{c}}, \mathbf{I})$ and $\rho \in (0,1)$ denotes the prescribed violation-probability limit. Following Zhang et al. (2025), we set $\mathbf{b} = \bar{\mathbf{c}} = \mathbf{1}$, $d = 1$, $n = 8$, and $\rho = 0.1$. The problem can be reformulated as second-order conic (SOC) problem, so we use the CVX (Grant et al., 2008) as the global reference. And we also compare our method with different types of SAA methods. We additionally evaluate GGDOpt First-order/Second-order (Zhang et al., 2025) and DiffOPT (Kong et al., 2025). For DiffOPT, we consistently use its iterative importance-sampling variant throughout our experiments, including settings with differentiable objectives, so that objective access remains value-only.

Results. On the linear Gaussian benchmark, Table 1 shows that the frozen constraint-only prior leaves substantial room for objective improvement. Without FKC, the sampler returns a mean objective value of -0.1172. Adding objective-value correction lowers it to -0.6519, only 0.0067 above the exact SOC optimum -0.6586. D^{3}Opt also attains the lowest mean and median and the smallest standard deviation 0.0015 among the evaluated diffusion methods. These gains are achieved at inference without objective gradients or objective-specific retraining.

Table 2: Objective-transfer results on ASQ and MW4 under the same linear Gaussian chance constraint. Each stochastic method is evaluated over 100 independent runs. Lower objective values and gaps are better.

Task	Method	FvalMean ↓	FvalStd ↓	FvalBest ↓	Optimality Gap (%) ↓
ASQ	SOC-CVX (Grant et al., 2008)	0.0979	–	0.0979	0.0000
	GGDOpt (Zhang et al., 2025) First-order ($\beta = 0.1$)	0.2026	0.0630	0.1055	106.9530
	GGDOpt (Zhang et al., 2025) Second-order ($\beta = 0.1, \sigma = 0.1$)	0.1994	0.0625	0.1055	103.6460
	DiffOPT (Kong et al., 2025)	0.1313	0.0081	0.1092	34.0980
	D^{3}Opt w/o FKC	0.8549	0.2086	0.3209	773.1010
	D^{3}Opt	**0.1035**	**0.0018**	**0.0999**	**5.7040**
MW4	SOC-CVX (Grant et al., 2008) (4-branch enumeration)	0.1762	–	0.1762	0.0000
	GGDOpt (Zhang et al., 2025) First-order ($\beta = 0.1$)	0.3132	0.0626	0.2168	77.7470
	GGDOpt (Zhang et al., 2025) Second-order ($\beta = 0.1, \sigma = 0.1$)	0.3123	0.0639	0.2159	77.2630
	DiffOPT (Kong et al., 2025)	0.2233	0.0132	0.1901	26.7250
	D^{3}Opt w/o FKC	0.7856	0.1518	0.4425	345.9200
	D^{3}Opt	**0.1834**	**0.0018**	**0.1798**	**4.0780**

5.2 Objective Generalization under a Linear Gaussian Chance Constraint

We next evaluate objective generalization by reusing the same frozen constraint-only diffusion checkpoint from Section 5.1 for objectives specified only after training. We consider a smooth anisotropic shifted quadratic (ASQ) objective and a non-convex, non-smooth four-well (MW4) objective, with respect to the same constraint as equation 11.

Anisotropic shifted quadratic (ASQ) objective. The ASQ objective can be described as:

$$f_{\mathrm{ASQ}}(\mathbf{x}) = \frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^\top \mathbf{Q}(\mathbf{x} - \boldsymbol{\mu}), \qquad \mathbf{Q} \succ 0 \tag{12}$$

The numerical parameters of Q and μ are provided in Appendix C.4. Since $\mathbf{Q} \succ 0$, the ASQ objective is smooth and strictly convex. Compared with the standard quadratic objective in Section 5.1, it changes both the unconstrained minimizer and the directional curvature while leaving the chance constraint unchanged. ASQ therefore provides a controlled test of whether the frozen constraint-only diffusion prior can adapt to a new smooth convex objective introduced only at inference time.

Non-convex and non-smooth four-well (MW4) objective. We further consider a four-well objective defined as the pointwise minimum of four positive-definite quadratic branches:

$$f_{\mathrm{MW4}}(\mathbf{x}) = \min_{m \in \{1,\dots,4\}} f_m(\mathbf{x}), \qquad f_m(\mathbf{x}) = \frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}_m)^\top \mathbf{Q}_m(\mathbf{x} - \boldsymbol{\mu}_m) + \alpha_m. \tag{13}$$

Each branch is strictly convex, whereas their pointwise minimum forms a multi-well non-convex objective with non-smooth branch-switching boundaries. Branch parameters are in Appendix C.4.

We compare against SOC-CVX, GGDOpt First-order and Second-order (Zhang et al., 2025), DiffOPT (Kong et al., 2025), and **D^{3}Opt** w/o FKC. We continue to use SOC-CVX to obtain the global optimal reference values. For MW4, the reference is obtained by solving the four convex quadratic branches separately and taking the minimum.

Results. Having learned the Gaussian chance-feasible structure once, D^{3}Opt applies the same frozen model to ASQ and MW4 using only inference-time objective evaluations. As shown in Table 2, FKC lowers mean objectives from 0.8549 to 0.1035 on ASQ and from 0.7856 to 0.1834 on MW4. The reported gaps to exact SOC references are 5.704% and 4.078%, versus DiffOPT's 34.098% and 26.725% under the reported protocol. D^{3}Opt also has the smallest standard deviation 0.0018 among learning-based methods on both tasks. These results demonstrate adaptation to smooth convex and non-smooth nonconvex objectives without retraining.

5.3 Chance-Constrained Economic Dispatch with Valve-Point Effects

We next evaluate D^{3}Opt on chance-constrained economic dispatch (CCED) with uncertain wind generation (Wang et al., 2017; Yang et al., 2021). Valve-point effects (VPE) introduce a non-smooth generation cost (Attaviriyanupap et al., 2002), providing a practical test of derivative-free optimiza-

Table 3: CCED-VPE-10 results at $\rho = 0.10$ over 100 runs, except PWL-MILP (one solve). Mean Probability is the average analytic satisfaction probability.

Method	Cost Mean ↓	Cost Std ↓	Cost Best ↓	Mean Probability
DE + exact projection	104399.055	5.2×10^{-8}	104399.055	0.9000
PWL-MILP + exact projection	104399.103	–	104399.103	0.9000
95% CI-SAA-DE	104400.092	≈ 0	104400.092	0.9003
GGDOpt (Zhang et al., 2025) (Vanilla)		Not directly applicable		
GGDOpt (Zhang et al., 2025) + VPE adaptation	108821.619	1389.720	106540.523	0.9522
DiffOPT (Kong et al., 2025)	105678.450	163.418	105240.143	0.9210
D^{3}Opt w/o FKC	107183.587	268.987	106146.715	0.9518
D^{3}Opt	**104823.577**	**102.065**	**104582.444**	0.9123

tion under chance constraints. The problem is formulated as:

$$\min_{\mathbf{p}} \quad C(\mathbf{p}) = \sum_{i=1}^{N_G} \left[a_i + b_i p_i + c_i p_i^2 + d_i \left| \sin\big(e_i (P_i^{\min} - p_i)\big) \right| \right], \tag{14}$$

$$\text{s.t.} \quad P_i^{\min} \le p_i \le P_i^{\max},\ i = 1, \ldots, N_G, \qquad \mathbb{P}_W \left\{ \sum_{i=1}^{N_G} p_i + W \ge P_D \right\} \ge 1 - \rho. \tag{15}$$

We consider a ten-unit system with $P_D = 2000$ MW, $W = 100Z$ MW, $Z \sim \mathrm{Beta}(5, 5)$, and $\rho = 0.10$. The generator parameters and additional evaluation details are provided in Appendix C.4.

We compare against DE (Balamurugan & Subramanian, 2008) and PWL-MILP (Pan et al., 2020), both with exact-feasibility projection, and a 95% CI-SAA-DE baseline combining confidence-bound SAA (Pagnoncelli et al., 2009) with DE. The projected DE solution serves as a numerical reference. Learning-based baselines include **D^{3}Opt** w/o FKC, GGDOpt (Zhang et al., 2025), and DiffOPT (Kong et al., 2025). Since GGDOpt requires derivatives, we evaluate a VPE first-order adaptation using branchwise derivatives, with $\mathrm{sign}(0) = 0$ at non-differentiable points.

Results. Table 3 shows that FKC reduces D^{3}Opt's mean cost from 107183.587 to 104823.577 (2.20%) and its cost standard deviation by 62.1% relative to the w/o FKC ablation. D^{3}Opt achieves the lowest mean and best costs and the smallest standard deviation among the evaluated learning-based methods; its mean cost is 0.81% below DiffOPT. Its mean analytic satisfaction probability is 0.9123, above the 0.90 requirement. Using only objective evaluations for the non-smooth valve-point cost, D^{3}Opt remains within 0.407% of the projected DE numerical reference.

6 Conclusion and Limitations

We presented a derivative-free diffusion framework for chance-constrained optimization that separates constraint-distribution learning from objective optimization. A risk-conditioned diffusion model is trained using constraint-only data and reused without retraining, while path-space FKC introduces post-specified objectives through function evaluations during reverse diffusion. Our analysis characterizes the terminal Boltzmann tilt, establishes feasibility preservation when the frozen prior is almost surely feasible, and separates optimization error into prior coverage, finite-particle approximation, and finite-temperature contributions. Experiments demonstrate objective transfer across smooth and non-smooth objectives under a fixed linear Gaussian chance constraint, with further results on valve-point economic dispatch, risk-level adaptation, and robust waveform design.

The method depends on the coverage of the learned prior: FK reweighting cannot reach regions outside its support, and relaxed empirical data acceptance does not guarantee exact feasibility. Our theoretical result bounds the raw best particle's expected objective gap; a screened or projected output requires separate analysis. The experimental comparisons use different inference budgets, and specialized solvers achieve lower costs on the analytically tractable CCED instance. Future work will improve prior coverage and particle efficiency and address higher-dimensional problems with less tractable feasibility checks.

AI USE STATEMENT

We used generative AI tools to assist with language polishing and the development and refinement of theoretical proofs. The authors reviewed and revised the AI-assisted text and checked the assumptions, logical steps, and conclusions of the proofs. We take responsibility for the final content of this work, including all text, theoretical claims, and proofs produced with the assistance of generative AI.

REPRODUCIBILITY STATEMENT

Constraint-only data generation and diffusion training are described in Section 4.1, while the annealed Feynman–Kac correction, temperature schedule, particle weighting, and resampling procedure are specified in Section 4.2. The problem instances, baseline settings, evaluation procedures, and main hyperparameters are reported in Sections 5.1, 5.2, and 5.3. Additional multi-risk, robust waveform, and ablation settings are provided in Appendices D.1, D.3, and D.2, respectively. The assumptions underlying the theoretical results are stated in the main text, with the Feynman–Kac derivation and complete proofs given in Appendices A, B.1, and B.2. Source code is available in the anonymous repository at `https://anonymous.4open.science/r/D3OPT/`.

REFERENCES

Pathom Attaviriyanupap, Hiroyuki Kita, Eiichi Tanaka, and Jun Hasegawa. A hybrid ep and sqp for dynamic economic dispatch with nonsmooth fuel cost function. *IEEE Transactions on Power Systems*, 17(2):411–416, 2002. doi: 10.1109/TPWRS.2002.1007911.

Charles Audet and John E. Dennis, Jr. Mesh adaptive direct search algorithms for constrained optimization. *SIAM Journal on Optimization*, 17(1):188–217, 2006. doi: 10.1137/040603371.

R. Balamurugan and S. Subramanian. Differential evolution-based dynamic economic dispatch of generating units with valve-point effects. *Electric Power Components and Systems*, 36(8):828–843, 2008. doi: 10.1080/15325000801911427.

Daniel Bienstock, Michael Chertkov, and Sean Harnett. Chance-constrained optimal power flow: Risk-aware network control under uncertainty. *Siam Review*, 56(3):461–495, 2014.

Yijia Cao, Yi Tan, Canbing Li, and Christian Rehtanz. Chance-constrained optimization-based unbalanced optimal power flow for radial distribution networks. *IEEE Transactions on Power Delivery*, 28(3):1855–1864, 2013. doi: 10.1109/TPWRD.2013.2259509.

Minshuo Chen, Song Mei, Jianqing Fan, and Mengdi Wang. Opportunities and challenges of diffusion models for generative ai. *National Science Review*, 11(12):nwae348, 2024. doi: 10.1093/nsr/nwae348.

Pierre Del Moral. *Feynman–Kac Formulae: Genealogical and Interacting Particle Systems with Applications*. Springer, 2004.

Shutong Ding, Yimiao Zhou, Ke Hu, Xi Yao, Junchi Yan, Xiaoying Tang, and Ye Shi. Exploring the boundary of diffusion-based methods for solving constrained optimization. *arXiv preprint arXiv:2502.10330*, 2025.

Thai Dinh, Ricardo Fukasawa, and James Luedtke. Exact algorithms for the chance-constrained vehicle routing problem. *Mathematical Programming*, 172(1):105–138, 2018.

Stewart N. Ethier and Thomas G. Kurtz. *Markov Processes: Characterization and Convergence*. Wiley, 1986.

Michael Grant, Stephen Boyd, and Yinyu Ye. CVX: Matlab software for disciplined convex programming, 2008.

Tawfik Guesmi, Anouar Farah, Ismail Marouani, Badr Alshammari, and Hsan Hadj Abdallah. Chaotic sine–cosine algorithm for chance-constrained economic emission dispatch problem including wind energy. *IET Renewable Power Generation*, 14(10):1808–1821, 2020. doi: 10.1049/iet-rpg.2019.1081.

Xiao-Li Hu, Thomas B. Schön, and Lennart Ljung. A general convergence result for particle filtering. *IEEE Transactions on Signal Processing*, 59(7):3424–3429, 2011. doi: 10.1109/TSP.2011.2135349.

Yeongmin Kim, Donghyeok Shin, Byeonghu Na, Minsang Park, Richard Lee Kim, and Il-Chul Moon. Lookahead sample reward guidance for test-time scaling of diffusion models. In *International Conference on Machine Learning*, 2026.

Lingkai Kong, Yuanqi Du, Wenhao Mu, Kirill Neklyudov, Valentin De Bortoli, Dongxia Wu, Haorui Wang, Aaron M. Ferber, Yian Ma, Carla P. Gomes, and Chao Zhang. Diffusion models as constrained samplers for optimization with unknown constraints. In *Proceedings of the 28th International Conference on Artificial Intelligence and Statistics*, volume 258 of *Proceedings of Machine Learning Research*, pp. 4582–4590. PMLR, 2025.

Tony Lelièvre, Mathias Rousset, and Gabriel Stoltz. *Free Energy Computations: A Mathematical Perspective*. Imperial College Press, 2010.

Xiner Li, Yulai Zhao, Chenyu Wang, Gabriele Scalia, Gokcen Eraslan, Surag Nair, Tommaso Biancalani, Shuiwang Ji, Aviv Regev, Sergey Levine, and Masatoshi Uehara. Derivative-free guidance in continuous and discrete diffusion models with soft value-based decoding. In *Advances in Neural Information Processing Systems*, volume 38, pp. 105884–105922, 2025.

Kazuyuki Masutomi, Yuichi Nagata, and Isao Ono. An evolutionary algorithm for black-box chance-constrained function optimization. *Journal of Advanced Computational Intelligence and Intelligent Informatics*, 17(2):272–282, 2013. doi: 10.20965/jaciii.2013.p0272.

Christian A. Naesseth, Fredrik Lindsten, and Thomas B. Schön. Elements of sequential monte carlo. *Foundations and Trends in Machine Learning*, 12(3), 2019.

Arkadi Nemirovski and Alexander Shapiro. Convex approximations of chance constrained programs. *SIAM Journal on Optimization*, 17(4):969–996, 2007. doi: 10.1137/050622328.

Zijing Ou, Chinmay Pani, and Yingzhen Li. Inference-time scaling of discrete diffusion models via importance weighting and optimal proposal design. In *International Conference on Learning Representations*, 2026.

Bernardo K. Pagnoncelli, Shabbir Ahmed, and Alexander Shapiro. Sample average approximation method for chance constrained programming: Theory and applications. *Journal of Optimization Theory and Applications*, 142(2):399–416, 2009. doi: 10.1007/s10957-009-9523-6.

Shanshan Pan, Jinbao Jian, Huangyue Chen, and Linfeng Yang. A full mixed-integer linear programming formulation for economic dispatch with valve-point effects, transmission loss and prohibited operating zones. *Electric Power Systems Research*, 180:106061, 2020. doi: 10.1016/j.epsr.2019.106061.

Xin-Yu Ren and Zhi-Hua Wang. Stochastic chance-constrained multi-objective optimization of interconnected multi-microgrids: Diverse energy storage configurations and interaction scenarios. *Energy*, pp. 141324, 2026.

Yinuo Ren, Wenhao Gao, Lexing Ying, Grant M. Rotskoff, and Jiequn Han. Driftlite: Lightweight drift control for inference-time scaling of diffusion models. In *International Conference on Learning Representations*, 2026.

Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pp. 234–241. Springer, 2015.

Raghav Singhal, Zachary Horvitz, Ryan Teehan, Mengye Ren, Zhou Yu, Kathleen McKeown, and Rajesh Ranganath. A general framework for inference-time scaling and steering of diffusion models. In *Proceedings of the 42nd International Conference on Machine Learning*, 2025.

Marta Skreta, Tara Akhound-Sadegh, Viktor Ohanesian, Roberto Bondesan, Alan Aspuru-Guzik, Arnaud Doucet, Rob Brekelmans, Alexander Tong, and Kirill Neklyudov. Feynman–kac correctors in diffusion: Annealing, guidance, and product of experts. In *Proceedings of the 42nd International Conference on Machine Learning*, 2025.

Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2021a.

Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021b. URL https://openreview.net/forum?id=PxTIG12RRHS.

Rainer Storn and Kenneth Price. Differential evolution – a simple and efficient heuristic for global optimization over continuous spaces. *Journal of Global Optimization*, 11(4):341–359, 1997. doi: 10.1023/A:1008202821328.

Kun-Yu Wang, Anthony Man-Cho So, Tsung-Hui Chang, Wing-Kin Ma, and Chong-Yung Chi. Outage constrained robust transmit optimization for multiuser miso downlinks: Tractable approximations by conic optimization. *IEEE Transactions on Signal Processing*, 62(21):5690–5705, 2014.

Peng Wang, Rujun Jiang, Qingyuan Kong, and Laura Balzano. A proximal DC algorithm for sample average approximation of chance constrained programming. *arXiv preprint arXiv:2301.00423*, 2023.

Zhiwen Wang, Chen Shen, Feng Liu, Xiangyu Wu, Cheng Liu, and Feng Gao. Chance-constrained economic dispatch with non-gaussian correlated wind power uncertainty. *IEEE Transactions on Power Systems*, 32(6):4880–4893, 2017. doi: 10.1109/TPWRS.2017.2672750.

Nick P. Whiteley. Sequential monte carlo samplers: Error bounds and insensitivity to initial conditions. *Stochastic Analysis and Applications*, 30(5):774–798, 2012. doi: 10.1080/07362994.2012.684323.

Yue Yang, Wenchuan Wu, Bin Wang, and Mingjie Li. Chance-constrained economic dispatch considering curtailment strategy of renewable energy. *IEEE Transactions on Power Systems*, 36(6):5792–5802, 2021. doi: 10.1109/TPWRS.2021.3081120.

Boyang Zhang, Zhiguo Wang, and Ya-Feng Liu. A gradient guided diffusion framework for chance constrained programming. In *Advances in Neural Information Processing Systems*, volume 38, 2025. doi: 10.52202/085713-4939.

Zeyang Zhang, Chuanhou Gao, and James Luedtke. New valid inequalities and formulations for the static joint chance-constrained lot-sizing problem: Z. zhang et al. *Mathematical Programming*, 199(1):639–669, 2023.

Shenglong Zhou, Lili Pan, Naihua Xiu, and Geoffrey Ye Li. A 0/1 constrained optimization solving sample average approximation for chance constrained programming. *Mathematics of Operations Research*, 2024. doi: 10.1287/moor.2023.0149.

A FEYNMAN–KAC CORRECTION

This appendix provides additional background on Feynman–Kac transforms and their connection to the weighted evolution of Markov processes. We first formulate the FK transform as a change of measure on trajectories and derive the associated backward and forward FK equations. We then recall the corresponding discrete-time FK model, which forms the basis of interacting-particle and sequential Monte Carlo approximations. Our presentation follows standard treatments of Feynman–Kac semigroups and interacting particle systems (Del Moral, 2004; Ethier & Kurtz, 1986; Lelièvre et al., 2010), as well as their recent use for inference-time correction and steering of diffusion models (Skreta et al., 2025; Singhal et al., 2025).

A.1 FROM THE FEYNMAN–KAC TRANSFORM TO THE FEYNMAN–KAC PDE

Feynman–Kac Transform. Let $(X_t)_{t\in[0,T]}$ be a diffusion on a state space $\mathcal{X} \subseteq \mathbb{R}^d$, with path law $\mathbb{Q}$,

$$dX_t = b_t(X_t)\,dt + \sigma_t(X_t)\,dB_t, \qquad X_0 \sim \mu_0,$$

where B_t is a standard Brownian motion. We write $a_t(x) := \sigma_t(x)\sigma_t(x)^\top$ and denote the infinitesimal generator of the process by

$$\mathcal{L}_t\varphi(x) = \langle b_t(x), \nabla\varphi(x)\rangle + \frac{1}{2}\,\mathrm{tr}\left(a_t(x)\nabla^2\varphi(x)\right)$$

for sufficiently regular test functions φ. Its adjoint $\mathcal{L}_t^*$ governs the marginal law q_t through the Fokker–Planck equation

$$\partial_t q_t = \mathcal{L}_t^* q_t.$$

Feynman–Kac transforms modify this baseline process through multiplicative trajectory weights (Del Moral, 2004; Lelièvre et al., 2010). Let $G : \mathcal{X} \to [0,\infty)$ be a terminal potential and let $V_t : \mathcal{X} \to \mathbb{R}$ be a measurable running potential. Assuming the corresponding exponential functional is integrable, define the transformed path law

$$\frac{d\mathbb{Q}^{G,V}}{d\mathbb{Q}} \triangleq \frac{1}{Z_{G,V}} \exp\left(-\int_0^T V_t(X_t)\,dt\right) G(X_T), \tag{16}$$

where

$$Z_{G,V} \triangleq \mathbb{E}_{\mathbb{Q}}\left[\exp\left(-\int_0^T V_t(X_t)\,dt\right) G(X_T)\right]$$

is the normalizing constant. Equation 16 makes explicit that the FK transform is fundamentally a change of measure on trajectories: paths can be favored through both a terminal criterion G and a running path criterion V_t.

To characterize this transformed path law, define the backward information function

$$h_t(x) \triangleq \mathbb{E}_{\mathbb{Q}}\left[\exp\left(-\int_t^T V_s(X_s)\,ds\right) G(X_T)\,\middle|\, X_t = x\right]. \tag{17}$$

The function $h_t(x)$ represents the expected future FK weight of a trajectory conditional on its current state. Under standard regularity conditions, the following proposition shows that it satisfies the backward Feynman–Kac equation.

Proposition 3 (Backward Feynman–Kac equation)**.** *Suppose the function h_t defined in equation 17 is sufficiently smooth and the standard integrability conditions required for Itô's formula hold. Then h_t satisfies*

$$\partial_t h_t(x) + \mathcal{L}_t h_t(x) - V_t(x)h_t(x) = 0, \qquad h_T(x) = G(x). \tag{18}$$

Proof. Fix $t \le u \le T$ and define

$$Z_{t,u} \triangleq \exp\left(-\int_t^u V_s(X_s)\,ds\right).$$

By the Markov property and the tower property,

$$h_t(x) = \mathbb{E}_{t,x}\left[Z_{t,u} h_u(X_u)\right]. \tag{19}$$

Applying Itô's formula to $Z_{t,u} h_u(X_u)$ yields

$$d(Z_{t,u} h_u(X_u)) = Z_{t,u}\left[\partial_u h_u + \mathcal{L}_u h_u - V_u h_u\right](X_u)\, du + Z_{t,u} \nabla h_u(X_u)^\top \sigma_u(X_u)\, dB_u.$$

In fact, the process

$$M_u \triangleq Z_{t,u} h_u(X_u), \qquad u \in [t, T],$$

is a martingale under the baseline law $\mathbb{Q}$. Indeed, for $t \le s \le u \le T$, using the multiplicative identity $Z_{t,u} = Z_{t,s} Z_{s,u}$, the Markov property, and the definition of h_s, we obtain

$$\mathbb{E}_{\mathbb{Q}}\left[M_u \mid \mathcal{F}_s\right] = Z_{t,s} \mathbb{E}_{\mathbb{Q}}\left[Z_{s,u} h_u(X_u) \mid \mathcal{F}_s\right] = Z_{t,s} h_s(X_s) = M_s.$$

Hence $\mathbb{E}_{t,x}[M_u] = M_t = h_t(x)$ is constant in u. The stochastic integral has zero conditional expectation, and hence the finite-variation term must vanish. Therefore,

$$\partial_u h_u + \mathcal{L}_u h_u - V_u h_u = 0.$$

The terminal condition follows immediately from equation 17 by setting $u = T$. □

The backward information function also provides a Markov representation of the transformed process. Let $P_{s,t}(x, dy)$ denote the transition kernel of the baseline process. The FK-transformed transition kernel can be written as

$$P^{h,V}_{s,t}(x, A) \triangleq \frac{1}{h_s(x)} \mathbb{E}_{s,x}\left[\exp\left(-\int_s^t V_r(X_r)\, dr\right) h_t(X_t) \mathbf{1}_{\{X_t \in A\}}\right]. \tag{20}$$

The backward FK identity guarantees that $P^{h,V}_{s,t}(x, \mathcal{X}) = 1$.

When h_t is sufficiently differentiable, the same transformed law can be represented by a Doob-type modification of the local drift. Its infinitesimal generator is

$$\mathcal{L}^h_t \varphi = \mathcal{L}_t \varphi + \langle a_t \nabla \log h_t, \nabla \varphi \rangle, \tag{21}$$

or equivalently,

$$dX_t = \left[b_t(X_t) + a_t(X_t) \nabla \log h_t(X_t)\right] dt + \sigma_t(X_t) dB^h_t.$$

This follows by expanding equation 20 over an infinitesimal time interval and using

$$\mathcal{L}_t(h\varphi) = h \mathcal{L}_t \varphi + \varphi \mathcal{L}_t h + \langle a_t \nabla h, \nabla \varphi \rangle$$

together with equation 18.

Remark 1 (Path-space transform versus drift realization)**.** *The drift correction $a_t \nabla \log h_t$ is one realization of the transformed path law when h_t is sufficiently regular; it is not the definition of the Feynman–Kac transform itself. The fundamental object is the change of measure in equation 16. For a prescribed marginal evolution, correction can in general be distributed differently among transport, diffusion, and reweighting terms (Skreta et al., 2025). This flexibility allows FK transforms to be implemented through weighted particle systems without requiring the transformed law to be represented solely through an additional gradient-based drift.*

For diffusion models, the baseline Markov process above can be taken to be the reverse diffusion sampler itself. For example, for the forward Ornstein–Uhlenbeck process

$$dY_\tau = -Y_\tau\, d\tau + \sqrt{2}\, dB_\tau,$$

the exact reverse process in generative time $t = T - \tau$ is

$$dX_t = \left[X_t + 2 \nabla \log q_{T-t}(X_t)\right] dt + \sqrt{2}\, d\widetilde{B}_t.$$

Here $a_t = 2 I_d$, and hence the corresponding Doob drift correction, when such a representation is used, is simply $2 \nabla \log h_t(x)$. This specialization recovers the familiar connection between diffusion guidance and Doob transforms, while Remark 1 highlights that the FK change of measure itself is more general.

Feynman–Kac PDE. A complementary characterization follows by evolving weighted marginal measures forward in time. Let $g_t : \mathcal{X} \to \mathbb{R}$ be a running log-weight rate and define the unnormalized FK measure through

$$\gamma_t(\varphi) \triangleq \mathbb{E}_{\mathbb{Q}}\left[\exp\left(\int_0^t g_s(X_s)\,ds\right)\varphi(X_t)\right] \tag{22}$$

for measurable test functions φ. Denote normalizing constant $Z_t := \mathbb{E}_{\mathbb{Q}}\left[\exp\left(\int_0^t g_s(X_s)\,ds\right)\right]$ and the normalized FK marginal by $p_t(\varphi) \triangleq \frac{\gamma_t(\varphi)}{Z_t}$.

Proposition 4 (Forward Feynman–Kac equation)**.** *Suppose the corresponding measures admit sufficiently regular densities. Then*

$$\partial_t \gamma_t = \mathcal{L}_t^* \gamma_t + g_t \gamma_t, \tag{23}$$

while the normalized FK marginal satisfies

$$\partial_t p_t = \mathcal{L}_t^* p_t + (g_t - p_t(g_t))\, p_t. \tag{24}$$

Proof. Let $W_t \triangleq \exp\left(\int_0^t g_s(X_s)\,ds\right)$. For a sufficiently regular test function φ, Itô's formula applied to $W_t\varphi(X_t)$ gives

$$\frac{d}{dt}\gamma_t(\varphi) = \gamma_t(\mathcal{L}_t\varphi) + \gamma_t(g_t\varphi).$$

Taking adjoints yields equation 23. Taking $\varphi \equiv 1$ gives

$$\dot{Z}_t = \gamma_t(g_t) = Z_t p_t(g_t).$$

Differentiating $p_t = \gamma_t / Z_t$ therefore yields

$$\partial_t p_t = \frac{\mathcal{L}_t^*\gamma_t + g_t\gamma_t}{Z_t} - \frac{\gamma_t}{Z_t^2}\dot{Z}_t = \mathcal{L}_t^* p_t + (g_t - p_t(g_t))\, p_t.$$

□

For the isotropic diffusion

$$dX_t = v_t(X_t)\,dt + \sigma_t\,dB_t,$$

equation 24 becomes

$$\partial_t p_t(x) = -\nabla\cdot(p_t(x)v_t(x)) + \frac{\sigma_t^2}{2}\Delta p_t(x) + (g_t(x) - p_t(g_t))\, p_t(x). \tag{25}$$

equation 25 makes explicit the three components of the FK evolution: transport, diffusion, and reweighting. This is the form used by Feynman–Kac Correctors for diffusion inference (Skreta et al., 2025).

Moreover, equation 22 immediately gives the weighted expectation identity

$$p_T(\varphi) = \frac{\mathbb{E}_{\mathbb{Q}}\left[W_T\varphi(X_T)\right]}{\mathbb{E}_{\mathbb{Q}}[W_T]}.$$

Thus FK marginals can be represented using multiplicatively weighted trajectories. This identity also motivates interacting-particle and sequential Monte Carlo approximations that redistribute particle mass during the evolution rather than retaining a single long-horizon importance weight (Del Moral, 2004; Naesseth et al., 2019; Skreta et al., 2025).

A.2 Discrete-Time Feynman–Kac Measures

The continuous formulation has a direct discrete-time analogue that is particularly convenient for diffusion samplers simulated on a finite time grid (Del Moral, 2004; Naesseth et al., 2019).

Consider $0 = t_0 < t_1 < \cdots < t_K = T$ and let

$$M_k(x_{k-1}, dx_k) \triangleq P_{t_{k-1},t_k}(x_{k-1}, dx_k)$$

denote the baseline Markov transition from t_{k-1} to t_k. Writing q_0 for the initial marginal, the baseline path law is

$$\mathbb{Q}_K(dx_{0:K}) = q_0(dx_0) \prod_{k=1}^{K} M_k(x_{k-1}, dx_k). \tag{26}$$

Let $G_k : \mathcal{X} \times \mathcal{X} \to [0, \infty)$ be measurable incremental potentials. The corresponding unnormalized discrete FK measure is defined, for a test function φ, by

$$\gamma_k(\varphi) \triangleq \mathbb{E}_{\mathbb{Q}} \left[\varphi(X_k) \prod_{j=1}^{k} G_j(X_{j-1}, X_j) \right]. \tag{27}$$

Its normalized marginal is

$$\eta_k(\varphi) \triangleq \frac{\gamma_k(\varphi)}{\mathbb{E}_{\mathbb{Q}} \left[\prod_{j=1}^{k} G_j(X_{j-1}, X_j) \right]}. \tag{28}$$

Equivalently, the associated unnormalized path measure on $(X_0, \ldots, X_k)$ is

$$\Gamma_k(dx_{0:k}) = q_0(dx_0) \prod_{j=1}^{k} \left[M_j(x_{j-1}, dx_j) G_j(x_{j-1}, x_j) \right].$$

Thus, as in continuous time, the FK construction preserves the baseline Markov dynamics as a reference process while modifying its law through multiplicative potentials.

The family $\{\eta_k\}_{k=0}^{K}$ constitutes a discrete Feynman–Kac flow. In general, evaluating these measures exactly is intractable, and interacting particle systems provide a standard Monte Carlo approximation (Del Moral, 2004). Sequential Monte Carlo maintains a finite population of particles, propagates them through the Markov transitions M_k, and uses the FK potentials to redistribute particle mass. Under standard assumptions, the resulting empirical measures consistently approximate the corresponding FK marginals as the number of particles tends to infinity (Naesseth et al., 2019).

Importantly, the choice of incremental potentials determines the sequence of intermediate FK measures and can therefore have a substantial effect on finite-particle efficiency, even when different constructions correspond to the same desired terminal change of measure. Recent diffusion-based FK methods exploit this flexibility to design intermediate potentials and proposal mechanisms for inference-time steering (Skreta et al., 2025; Singhal et al., 2025).

B Proofs

B.1 Feasibility Preservation under Feynman–Kac Tilting

We provide the proof of Proposition 1 for the discrete-time Feynman–Kac construction used in Section 4.2.

Proof of Proposition 1. Fix a risk level ρ and consider the path space $\Omega_K = \mathcal{X}^{K+1}$, $X_{0:K} = (X_0, \ldots, X_K)$. The path law of the frozen conditional reverse diffusion is

$$\mathbb{Q}_\theta(dx_{0:K} \mid \rho) = p_{\mathrm{ref}}(dx_0) \prod_{k=1}^{K} M_k^\theta(x_{k-1}, dx_k \mid \rho).$$

Let $L_K(x_{0:K}; \rho) \triangleq \prod_{k=1}^{K} G_k(x_{k-1}, x_k; \rho) \geq 0$ denote the Feynman–Kac path weight, and assume that $0 < Z_K(\rho) \triangleq \mathbb{E}_{Q_\theta(\cdot|\rho)} [L_K(X_{0:K}; \rho)] < \infty$. The normalized Feynman–Kac path measure is therefore

$$\mathbb{Q}_\theta^{\mathrm{FK}}(dx_{0:K} \mid \rho) = \frac{L_K(x_{0:K}; \rho)}{Z_K(\rho)} \mathbb{Q}_\theta(dx_{0:K} \mid \rho) \Rightarrow \frac{d\mathbb{Q}_\theta^{\mathrm{FK}}(\cdot \mid \rho)}{d\mathbb{Q}_\theta(\cdot \mid \rho)}(X_{0:K}) = \frac{L_K(X_{0:K}; \rho)}{Z_K(\rho)}.$$

Since the right-hand side is a well-defined nonnegative density,

$$\mathbb{Q}_\theta^{\mathrm{FK}}(\cdot \mid \rho) \ll \mathbb{Q}_\theta(\cdot \mid \rho). \tag{29}$$

We next transfer this absolute continuity from path space to the terminal marginal. Define the measurable terminal evaluation map $e_K : \Omega_K \to \mathcal{X}$, $e_K(x_{0:K}) = x_K$. By definition,

$$p_\theta(\cdot \mid \rho) = (e_K)_\# \mathbb{Q}_\theta(\cdot \mid \rho), \qquad \pi_{\theta,\beta_K}(\cdot \mid \rho) = (e_K)_\# \mathbb{Q}_\theta^{\mathrm{FK}}(\cdot \mid \rho). \tag{30}$$

where $(e_K)_\#$ denotes the pushforward under the terminal evaluation map $e_K(x_{0:K}) = x_K$.

Let $A \subseteq \mathcal{X}$ be any measurable set satisfying $p_\theta(A \mid \rho) = 0$. Using the first identity in equation 30, $\mathbb{Q}_\theta\left(e_K^{-1}(A) \mid \rho\right) = 0$. The path-space absolute continuity therefore implies $\mathbb{Q}_\theta^{\mathrm{FK}}\left(e_K^{-1}(A) \mid \rho\right) = 0$. Applying the second identity in equation 30 gives $\pi_{\theta,\beta_K}(A \mid \rho) = 0$. Since this holds for every measurable A with $p_\theta(A \mid \rho) = 0$, we conclude that

$$\pi_{\theta,\beta_K}(\cdot \mid \rho) \ll p_\theta(\cdot \mid \rho). \tag{31}$$

Finally, suppose that the frozen diffusion is chance-feasible almost surely, $p_\theta(\mathcal{X}_\rho \mid \rho) = 1$. Equivalently, $p_\theta(\mathcal{X}_\rho^c \mid \rho) = 0$. Taking $A = \mathcal{X}_\rho^c$ in equation 31 yields $\pi_{\theta,\beta_K}(\mathcal{X}_\rho^c \mid \rho) = 0$ and therefore

$$\pi_{\theta,\beta_K}(\mathcal{X}_\rho \mid \rho) = 1.$$

Thus, any almost-sure chance feasibility encoded by the frozen diffusion prior is preserved under Feynman–Kac correction. □

For the objective-only construction in Section 4.2, the telescoping identity gives

$$L_K(X_{0:K}; \rho) = \prod_{k=1}^{K} G_k(X_{k-1}, X_k; \rho) = \exp[-\beta_K f(X_K)].$$

Hence, whenever $0 < Z_{\theta,\beta_K}(\rho) = \int_{\mathcal{X}} e^{-\beta_K f(x)} p_\theta(dx \mid \rho) < \infty$, the terminal Radon–Nikodym derivative is explicitly

$$\frac{d\pi_{\theta,\beta_K}(\cdot \mid \rho)}{dp_\theta(\cdot \mid \rho)}(x) = \frac{e^{-\beta_K f(x)}}{Z_{\theta,\beta_K}(\rho)}. \tag{32}$$

This makes the feasibility-preservation property explicit for the proposed objective-only Feynman–Kac correction.

B.2 Proof of Proposition 2

Setup and error decomposition. Assumption 1 guarantees the fixed-horizon particle approximation bound required below. Throughout the proof, we fix ρ and omit conditioning on ρ whenever no ambiguity arises. We denote the frozen reverse-diffusion kernels by $M_k = M_k^\theta(\cdot, \cdot \mid \rho)$ and the incremental potentials by G_k. By the telescoping construction in Section 4.2, the exact terminal FK marginal is

$$\eta_K = \pi_{\theta,\beta_K}, \qquad \pi_{\theta,\beta_K}(dx) = \frac{e^{-\beta_K f(x)} p_\theta(dx)}{Z_{\theta,\beta_K}}.$$

Let η_K^M denote its M-particle SMC approximation. Adding and subtracting $\pi_{\theta,\beta_K}(f)$ and $f^\star_{\theta,\rho}$ gives

$$\left|\mathbb{E}[\eta_K^M(f)] - f_\rho^\star\right| \le \underbrace{\left|\mathbb{E}[\eta_K^M(f)] - \pi_{\theta,\beta_K}(f)\right|}_{I_M} + \underbrace{\left|\pi_{\theta,\beta_K}(f) - f^\star_{\theta,\rho}\right|}_{I_\beta} + \underbrace{\left|f^\star_{\theta,\rho} - f_\rho^\star\right|}_{I_{\mathrm{cov}}}. \tag{33}$$

We bound I_M, I_β, and I_{cov} separately.

Step 1: Finite inverse-temperature error. We first bound I_β. By definition,

$$f^\star_{\theta,\rho} = \operatorname*{ess\,inf}_{x\sim p_\theta} f(x),$$

and, for every $\epsilon > 0$,

$$a_{\theta,\epsilon} = p_\theta\left(f(X) \le f^\star_{\theta,\rho} + \epsilon\right) > 0.$$

Hence the normalizing constant of the tilted distribution satisfies

$$Z_{\theta,\beta_K} = \int_{\mathcal{X}} e^{-\beta_K f(x)} p_\theta(dx) \ge \int_{\{f \le f^\star_{\theta,\rho}+\epsilon\}} e^{-\beta_K f(x)} p_\theta(dx) \ge a_{\theta,\epsilon} \exp\left[-\beta_K\left(f^\star_{\theta,\rho} + \epsilon\right)\right]. \tag{34}$$

Moreover, $\frac{d\pi_{\theta,\beta_K}}{dp_\theta}(x) = \frac{e^{-\beta_K f(x)}}{Z_{\theta,\beta_K}}$, and therefore

$$D_{\mathrm{KL}}\left(\pi_{\theta,\beta_K} \,\|\, p_\theta\right) = -\beta_K \pi_{\theta,\beta_K}(f) - \log Z_{\theta,\beta_K}.$$

Since relative entropy is nonnegative,

$$\pi_{\theta,\beta_K}(f) \le -\frac{1}{\beta_K}\log Z_{\theta,\beta_K}.$$

Combining this inequality with equation 34 yields, for every $\epsilon > 0$,

$$\pi_{\theta,\beta_K}(f) - f^\star_{\theta,\rho} \le \epsilon + \frac{1}{\beta_K}\log\frac{1}{a_{\theta,\epsilon}}. \tag{35}$$

Proposition 1 gives $\pi_{\theta,\beta_K} \ll p_\theta$. Since

$$f(x) \ge f^\star_{\theta,\rho} \qquad p_\theta\text{-a.s.},$$

the same inequality holds π_{θ,β_K}-almost surely. Consequently,

$$\left|\pi_{\theta,\beta_K}(f) - f^\star_{\theta,\rho}\right| = \pi_{\theta,\beta_K}(f) - f^\star_{\theta,\rho}.$$

Taking the infimum of equation 35 over $\epsilon > 0$ gives

$$I_\beta \le \inf_{\epsilon>0}\left\{\epsilon + \frac{1}{\beta_K}\log\frac{1}{a_{\theta,\epsilon}}\right\}. \tag{36}$$

The right-hand side vanishes as $\beta_K \to \infty$. Indeed, for any $\delta > 0$, choose $\epsilon \in (0, \delta)$. Since ϵ is fixed and $a_{\theta,\epsilon} > 0$, for all sufficiently large β_K,

$$\frac{1}{\beta_K}\log\frac{1}{a_{\theta,\epsilon}} \le \delta - \epsilon.$$

Thus

$$\lim_{\beta_K\to\infty}\inf_{\epsilon>0}\left\{\epsilon + \frac{1}{\beta_K}\log\frac{1}{a_{\theta,\epsilon}}\right\} = 0. \tag{37}$$

Step 2: Finite-particle Feynman–Kac error. By Assumption 1,

$$\left(\mathbb{E}\left[\left|\eta^M_K(f) - \pi_{\theta,\beta_K}(f)\right|^2\right]\right)^{1/2} \le \frac{C_{K,f}}{\sqrt{M}}.$$

Consequently, by Cauchy–Schwarz,

$$\begin{aligned} I_M &= \left|\mathbb{E}[\eta^M_K(f)] - \pi_{\theta,\beta_K}(f)\right| \le \mathbb{E}\left[\left|\eta^M_K(f) - \pi_{\theta,\beta_K}(f)\right|\right] \\ &\le \left(\mathbb{E}\left[\left|\eta^M_K(f) - \pi_{\theta,\beta_K}(f)\right|^2\right]\right)^{1/2} \le \frac{C_{K,f}}{\sqrt{M}}. \end{aligned} \tag{38}$$

Step 3: Coverage error and conclusion. By Assumption 1, $p_\theta(\mathcal{X}_\rho \mid \rho) = 1$. Since $f(x) \geq f^\star_\rho$, $x \in \mathcal{X}_\rho$, we have

$$f(x) \geq f^\star_\rho \qquad p_\theta\text{-a.s.}$$

and hence $f^\star_{\theta,\rho} \geq f^\star_\rho$. Therefore,

$$I_{\text{cov}} = \left| f^\star_{\theta,\rho} - f^\star_\rho \right| = f^\star_{\theta,\rho} - f^\star_\rho = \Delta_{\theta,\rho}. \tag{39}$$

Substituting equation 36, equation 38, and equation 39 into the decomposition equation 33 gives

$$\left| \mathbb{E}[\eta^M_K(f)] - f^\star_\rho \right| \leq \Delta_{\theta,\rho} + \frac{C_{K,f}}{\sqrt{M}} + \inf_{\epsilon>0} \left\{ \epsilon + \frac{1}{\beta_K} \log \frac{1}{a_{\theta,\epsilon}} \right\}. \tag{40}$$

Finally, for each fixed finite inverse-temperature schedule, $C_{K,f} < \infty$ is independent of M. Thus, letting $M \to \infty$ first removes the finite-particle term. Subsequently, letting $\beta_K \to \infty$ and using equation 37 yields

$$\limsup_{\beta_K \to \infty} \limsup_{M \to \infty} \left| \mathbb{E}[\eta^M_K(f)] - f^\star_\rho \right| \leq \Delta_{\theta,\rho}, \tag{41}$$

which establishes equation 10.

C Implementation Details and Hyper-parameters

C.1 Constraint-Only Data Generation

Training decisions are generated and filtered using only constraint information. Each dataset contains 10,000 retained samples. For every configuration, training risk levels are sampled uniformly from $[0, 0.5]$; the risk levels reported in the experiments specify evaluation conditions.

Linear Gaussian feasibility is evaluated analytically. CCED uses two independent sets of $100{,}000$ wind scenarios for labeling and validation. Waveform data generation uses $1{,}000$ preliminary filtering scenarios, $4{,}000$ labeling scenarios, and $5{,}000$ validation scenarios without acceptance relaxation.

C.2 Diffusion Training

We use conditional 1D U-Nets for the linear Gaussian and robust waveform experiments and a conditional 2D U-Net for CCED-VPE-10. The diffusion models use $1{,}000$ forward diffusion steps with a linear noise-variance schedule from 10^{-4} to 0.02. Training uses Adam with learning rate 10^{-4} and batch size 256. The models are frozen before downstream optimization. ASQ and MW4 reuse the same frozen model as the original linear Gaussian experiment, without additional data generation or training. A single CCED model is likewise reused across all evaluated risk levels.

C.3 Feynman–Kac Inference

Post-specified objectives are introduced through derivative-free Feynman–Kac correction using only objective-function evaluations. We use 512 particles for all inference experiments. The prescribed inverse-temperature schedule follows equation 6 and is synchronized with the reverse-diffusion steps, with β_k specifying the objective tilt at step k. The schedule is fixed before sampling, while ESS is used only to determine when particle resampling is performed.

C.4 Problem-Specific Settings

Linear Gaussian and objective transfer. For the linear Gaussian benchmark, we use $\mathbf{b} = \bar{\mathbf{c}} = \mathbf{1}$, $d = 1$, $n = 8$, and $\rho = 0.1$, following (Zhang et al., 2025). The linear Gaussian, ASQ, and MW4 experiments share the same chance constraint, whose feasibility is evaluated through its exact SOC reformulation. SOC-CVX provides the global reference.

Transfer experiment to ASQ. The parameter in ASQ problem equation 12, we set

$$\mathbf{Q} = \mathrm{diag}(0.6, 0.8, 1.0, 1.2, 1.4, 1.6, 1.8, 2.0),$$

and

$$\boldsymbol{\mu} = (-0.55, -0.45, -0.35, -0.25, -0.15, -0.05, 0.05, 0.15)^\top.$$

Transfer experiment to MW4. In MW4 problem equation 13, the four quadratic branches use $\boldsymbol{\mu}_m = -0.25\mathbf{1} + 0.22\mathbf{s}_m$, $(\alpha_1, \alpha_2, \alpha_3, \alpha_4) = (0.025, 0.015, 0, 0.035)$, and $\mathbf{Q}_m = \mathrm{diag}(\mathrm{roll}_{2(m-1)}(\mathbf{q}))$, where $q_j = 0.8 + 0.4(j-1)/7$. The sign vectors $\mathbf{s}_m$ are specified by

$$\mathbf{S} = \begin{bmatrix} 1 & 1 & 1 & 1 & -1 & -1 & -1 & -1 \\ 1 & 1 & -1 & -1 & 1 & 1 & -1 & -1 \\ 1 & -1 & 1 & -1 & 1 & -1 & 1 & -1 \\ 1 & 1 & -1 & -1 & -1 & -1 & 1 & 1 \end{bmatrix}.$$

For MW4, the global reference is obtained by solving the four convex quadratic branches separately under the shared chance constraint and taking the smallest optimum. Reference solutions are used only for post-hoc evaluation and are never supplied to the sampling methods.

Table 4: Generator output limits and valve-point cost coefficients for CCED-VPE-10.

Unit i	$P_i^{\min}(MW)$	$P_i^{\max}(MW)$	a_i	b_i	c_i	d_i	e_i
1	10	55	1000.403	40.5407	0.12951	33	0.0174
2	20	80	950.606	39.5804	0.10908	25	0.0178
3	47	120	900.705	36.5104	0.12511	32	0.0162
4	20	130	800.705	39.5104	0.12111	30	0.0168
5	50	160	756.799	38.5390	0.15247	30	0.0148
6	70	240	451.325	46.1592	0.10587	20	0.0163
7	60	300	1243.531	38.3055	0.03546	20	0.0152
8	70	340	1049.998	40.3965	0.02803	30	0.0128
9	135	470	1658.569	36.3278	0.02111	60	0.0136
10	150	470	1356.659	38.2704	0.01799	40	0.0141

CCED-VPE-10. We use a ten-unit system with $P_D = 2000$ MW, $W = 100Z$ MW, $Z \sim \mathrm{Beta}(5, 5)$, and $\rho = 0.10$ in the main experiment. Table 4 gives the generator limits and valve-point cost coefficients.

Under this uncertainty model, the chance constraint is equivalent to

$$\sum_{i=1}^{N_G} p_i \geq P_D - 100F^{-1}_{\mathrm{Beta}(5,5)}(\rho). \tag{42}$$

For CCED, physical outputs are normalized as $x_i = 2(p_i - P_i^{\min})/(P_i^{\max} - P_i^{\min}) - 1$ and padded with zeros to form a 16-dimensional input. The inference objective is scaled as $\widetilde{C}(\mathbf{p}) = (C(\mathbf{p}) - 100000)/10000$. Reported costs always use the original, unscaled objective.

C.5 Evaluation

For each returned CCED decision $\mathbf{p}^{(r)}$, its satisfaction probability is evaluated analytically as

$$P_r = 1 - F_{\mathrm{Beta}(5,5)}\left(\frac{P_D - \sum_{i=1}^{N_G} p_i^{(r)}}{100}\right). \tag{43}$$

The cumulative distribution function is set to zero below its support and one above it. Mean Probability in Table 3 is the average of P_r across independent runs; for PWL-MILP, it is the probability of its single solution. Costs and probabilities are evaluated on the same returned decisions. Generator bounds and $P_r \geq 1 - \rho$ are checked for every run with a numerical tolerance of 10^{-10}. Higher satisfaction probabilities indicate larger reliability margins and do not necessarily imply lower cost.

Reported statistics. Each stochastic method is evaluated over 100 independent runs per setting. We compute the reported mean, sample standard deviation, and, where applicable, median from the decisions returned by the 100 runs. "Best" is the lowest cost among these returned decisions. PWL-MILP is solved once per risk level and therefore has no reported standard deviation.

D Additional Experimental Results

Table 5: Risk-level adaptation on CCED-VPE-10 at $\rho \in \{0.05, 0.15, 0.20\}$. We report mean $\pm$ sample standard deviation and best cost over 100 runs, except PWL-MILP, which uses one solve per risk level. Bold costs indicate the best available learning-based results.

ρ	Method	Cost Mean ↓	Cost Std ↓	Cost Best ↓	Mean Probability
0.05	DE + exact projection	104689.448	0.000	104689.448	0.9500
	PWL-MILP + exact projection	104689.483	–	104689.483	0.9500
	95% CI-SAA-DE	104717.493	0.000	104717.493	0.9537
	GGDOpt (Zhang et al., 2025) + VPE adaptation	108951.188	1344.889	106546.719	0.9729
	DiffOPT (Kong et al., 2025)	105986.707	166.335	105554.119	0.9500
	D^{3}Opt	**105111.845**	**112.923**	**104831.074**	0.9500
0.15	DE + exact projection	104190.322	0.000	104190.322	0.8500
	PWL-MILP + exact projection	104190.504	–	104190.504	0.8500
	95% CI-SAA-DE	104197.316	0.000	104197.316	0.8519
	GGDOpt (Zhang et al., 2025) + VPE adaptation	108734.612	1421.301	106437.094	0.9331
	DiffOPT (Kong et al., 2025)	105455.728	162.442	105014.749	0.8500
	D^{3}Opt	**104590.257**	**113.299**	**104284.908**	0.8500
0.20	DE + exact projection	104018.596	0.000	104018.596	0.8000
	PWL-MILP + exact projection	104018.932	–	104018.932	0.8000
	95% CI-SAA-DE	104020.643	0.000	104020.643	0.8007
	GGDOpt (Zhang et al., 2025) + VPE adaptation	108670.804	1445.010	106261.099	0.9166
	DiffOPT (Kong et al., 2025)	105271.029	159.120	104860.692	0.8000
	D^{3}Opt	**104449.668**	**110.163**	**104166.527**	0.8000

D.1 Multi-Risk Generalization on CCED-VPE-10

We investigate whether the frozen constraint-only diffusion model used in Section 5.3 can support different risk requirements without retraining. We retain the same dispatch objective, generator parameters, and Beta wind distribution, and evaluate $\rho \in \{0.05, 0.10, 0.15, 0.20\}$. These test levels lie within the continuous risk interval used for training. At inference time, we change only the risk condition supplied to the frozen model, while feasibility is evaluated under the corresponding chance constraint.

Baselines and Evaluation We compare against DE with exact-feasibility projection and DiffOPT (Kong et al., 2025). SAA-DE is solved separately for each risk level, while DiffOPT and our method reuse the same frozen checkpoint across all four levels.

Each stochastic method is evaluated over 100 independent runs per risk level. We report the mean, sample standard deviation, and best cost under the original valve-point objective, together with the number of returned solutions satisfying the analytic chance constraint under the Beta wind distribution.

Results Table 5 shows that our method achieves lower mean and best costs than DiffOPT at every tested risk level, together with smaller cost standard deviations. All 100 returned solutions at each risk level satisfy the analytic chance constraint. As ρ increases, the required thermal generation decreases, and the costs decrease accordingly.

These results demonstrate that the same frozen constraint-only model can support optimization under different risk requirements within the training interval without data regeneration or retraining. The traditional methods retain lower costs on this analytically tractable instance, while the comparison with DiffOPT supports the effectiveness of derivative-free Feynman–Kac correction under the shared checkpoint and evaluation protocol.

Table 6: Ablation experiments on CCED-VPE-10 at $\rho = 0.1$, varying particle count $M \in \{64, 256, 512, 1024, 2048\}$ and terminal inverse temperature $\beta_{\max} \in \{0, 5, 10, 20, 50\}$. Results are reported over 100 runs per setting.

$\beta_{\max}$	M	FvalMean ↓	FvalStd ↓	FvalBest ↓
10	64	105161.669	243.555	104691.425
10	256	104887.677	126.620	104593.105
10	512	104823.577	102.065	104582.444
10	1024	104759.778	90.430	104549.575
10	2048	**104707.598**	**74.224**	**104521.263**
0	512	107183.587	268.987	106146.715
5	512	104977.214	124.773	104665.561
10	512	104823.577	**102.065**	104582.444
20	512	104768.454	104.858	104523.105
50	512	**104759.376**	103.180	**104519.511**

D.2 Ablation Study

We examine the effects of particle count M and terminal inverse temperature $\beta_{\max}$ on CCED-VPE-10 at $\rho = 0.1$. In each ablation, we vary one parameter while keeping the other and all remaining settings fixed. The setting $\beta_{\max} = 0$ disables objective tilting during reverse diffusion and therefore serves as an unguided baseline. Each configuration is evaluated over 100 independent runs.

Table 6 shows that increasing the particle count consistently improves both solution quality and stability. As M increases from 64 to 2048 with $\beta_{\max} = 10$, the mean cost decreases from 105161.669 to 104707.598, while the sample standard deviation decreases from 243.555 to 74.224. The best cost also improves from 104691.425 to 104521.263. These results indicate that a larger particle population improves coverage of low-objective regions and reduces finite-particle variability.

Increasing $\beta_{\max}$ also strengthens objective optimization. With $M = 512$, raising $\beta_{\max}$ from 0 to 50 reduces the mean cost from 107183.587 to 104759.376 and the best cost from 106146.715 to 104519.511. The standard deviation decreases substantially up to $\beta_{\max} = 10$ and changes only slightly at $\beta_{\max} = 20, 50$. This trend shows that stronger objective tilting moves the particle system toward lower-cost regions, while overly strong concentration may trade stability for additional objective improvement.

The observed trends are consistent with Proposition 2: increasing M reduces the finite-particle contribution, while increasing $\beta_{\max}$ reduces the finite-temperature contribution. We use $M = 512$ and $\beta_{\max} = 10$ as the default configuration because it provides a favorable balance between computational cost, solution quality, and stability. Larger settings further improve the mean or best cost but require more particles or produce slightly higher variability.

D.3 Robust Waveform Experiments

We further evaluate the method on robust multi-user waveform design with $N_t = 16$ transmit antennas and $K = 3$ users. Let $\mathbf{w}_k \in \mathbb{R}^{N_t}$ denote the beamforming vector for user k, and let $\mathbf{W} = (\mathbf{w}_1, \ldots, \mathbf{w}_K)$. The problem minimizes total transmission power subject to individual signal-to-interference-plus-noise ratio chance constraints:

$$\begin{aligned} \min_{\mathbf{W}} \quad & f(\mathbf{W}) = \sum_{k=1}^{K} \|\mathbf{w}_k\|_2^2, \\ \text{s.t.} \quad & \mathbb{P}\left\{ \frac{(\mathbf{h}_k^\top \mathbf{w}_k)^2}{\sum_{j \neq k} (\mathbf{h}_k^\top \mathbf{w}_j)^2 + \sigma_i^2} \geq \gamma \right\} \geq 1 - \rho, \qquad k = 1, \ldots, K. \end{aligned} \tag{44}$$

The uncertain channel is modeled as $\mathbf{h}_k = \bar{\mathbf{h}}_k + \mathbf{e}_k$, where $\mathbf{e}_k \sim \mathcal{N}(\mathbf{0}, \sigma_e^2 \mathbf{I})$. Following GGDOpt (Zhang et al., 2025), we set $\gamma = 0.5$, $\sigma_e^2 = 0.002$, and $\sigma_i = 1$, and evaluate $\rho \in \{0.05, 0.10, 0.15, 0.20\}$.

Table 7: Robust waveform results with $N_t = 16$ and $K = 3$. Learning-based power statistics are from 100 runs. The final column reports per-user chance-constraint satisfaction probabilities on the returned decisions, averaged across runs for learning-based methods. For a traditional baseline it refers to its single solution. Bold power values are the lowest mean and best among learning-based methods.

ρ	Method	Power Mean ↓	Power Std ↓	Power Best ↓	$P_1/P_2/P_3$
0.05	Sphere Bounding (Wang et al., 2014)	0.1374	–	–	(0.99; 0.99; 0.99)
	Bernstein-type Inequality (Wang et al., 2014)	0.1260	–	–	(0.97; 0.98; 0.95)
	GGDOpt First-order (Zhang et al., 2025)	0.127992	0.000205	0.127339	(0.99; 0.98; 0.95)
	GGDOpt Second-order (Zhang et al., 2025)	0.127989	0.000206	0.127335	(0.98; 0.98; 0.95)
	DiffOPT (Kong et al., 2025)	0.246433	0.000565	0.246110	(0.99; 0.98; 0.97)
	D^{3}Opt	**0.127732**	0.000209	**0.127275**	(0.97; 0.95; 0.96)
0.10	Sphere Bounding (Wang et al., 2014)	0.1366	–	–	(0.99; 0.99; 0.99)
	Bernstein-type Inequality (Wang et al., 2014)	0.1253	–	–	(0.93; 0.94; 0.93)
	GGDOpt First-order (Zhang et al., 2025)	0.127941	0.000251	0.127019	(0.92; 0.91; 0.97)
	GGDOpt Second-order (Zhang et al., 2025)	0.127938	0.000252	0.127009	(0.95; 0.92; 0.92)
	DiffOPT (Kong et al., 2025)	0.237796	0.000598	0.236225	(0.97; 0.94; 0.94)
	D^{3}Opt	**0.126043**	0.000187	**0.125518**	(0.92; 0.90; 0.91)
0.15	Sphere Bounding (Wang et al., 2014)	0.1361	–	–	(0.99; 0.99; 0.99)
	Bernstein-type Inequality (Wang et al., 2014)	0.1248	–	–	(0.90; 0.90; 0.92)
	GGDOpt First-order (Zhang et al., 2025)	0.127983	0.000187	0.127474	(0.90; 0.87; 0.92)
	GGDOpt Second-order (Zhang et al., 2025)	0.127981	0.000189	0.127467	(0.90; 0.87; 0.86)
	DiffOPT (Kong et al., 2025)	0.245401	0.000475	0.245118	(0.90; 0.92; 0.89)
	D^{3}Opt	**0.125047**	0.000177	**0.124726**	(0.85; 0.87; 0.86)
0.20	Sphere Bounding (Wang et al., 2014)	0.1357	–	–	(0.99; 0.99; 0.99)
	Bernstein-type Inequality (Wang et al., 2014)	0.1244	–	–	(0.90; 0.89; 0.91)
	GGDOpt First-order (Zhang et al., 2025)	0.127964	0.000245	0.127006	(0.90; 0.82; 0.88)
	GGDOpt Second-order (Zhang et al., 2025)	0.127962	0.000247	0.126994	(0.81; 0.85; 0.88)
	DiffOPT (Kong et al., 2025)	0.211671	0.000463	0.210839	(0.88; 0.87; 0.91)
	D^{3}Opt	**0.124361**	0.000184	**0.123822**	(0.85; 0.82; 0.80)

Baselines and Evaluation We compare against Sphere Bounding (Wang et al., 2014), Bernstein-type Inequality (Wang et al., 2014), first- and second-order GGDOpt (Zhang et al., 2025), and DiffOPT (Kong et al., 2025). DiffOPT uses iterative importance sampling, whereas D^{3}Opt applies objective-value-based Feynman–Kac correction during reverse diffusion.

Each learning-based method is evaluated over 100 runs. We report mean, sample standard deviation, and best power. For user k, let $\hat{P}_{r,k}$ be the fraction of test scenarios satisfying its SINR constraint for returned decision r. Table 7 reports $\bar{P}_k = 100^{-1} \sum_{r=1}^{100} \hat{P}_{r,k}$ for learning-based methods; traditional-baseline triplets refer to one solution. These run averages do not give the number of individually feasible returns.

Results Table 7 shows that D^{3}Opt achieves the lowest mean and best transmission power among the evaluated learning-based methods at all four risk levels. The minimum of its three run-averaged user satisfaction probabilities is 0.95, 0.90, 0.85, and 0.80 for $\rho = 0.05$, 0.10, 0.15, and 0.20, respectively, matching the nominal $1 - \rho$ thresholds at the reported two-decimal precision. Relative to second-order GGDOpt, D^{3}Opt reduces mean power by 0.20%, 1.48%, 2.29%, and 2.81%. For D^{3}Opt, mean power decreases monotonically as ρ increases. These results are consistent with using the allowed risk budget closely across the tested risk levels without retraining.