GOAL-DRIVEN BAYESIAN OPTIMAL EXPERIMENTAL DESIGN FOR ROBUST DECISION-MAKING UNDER MODEL UNCERTAINTY

Anonymous authors
Paper under double-blind review

ABSTRACT

Bayesian optimal experimental design (BOED) aims to predict experiments that can optimally reduce the uncertainty in the model parameters. However, in many decision-critical applications, accurate parameter estimation does not necessarily translate to better decision-making, as not all parameters may significantly affect the efficacy of the decisions made in the presence of uncertainty. In this work, we propose GoBOED (Goal-driven Bayesian Optimal Experimental Design) to directly optimize the experimental design to reduce the model uncertainty that critically affects the quality of the downstream decision-making task of interest. We establish a computationally tractable connection between BOED and robust optimal control based on an uncertain model through convex optimization. This new integrated framework for robust control under uncertainty enables efficient gradient computation through a decision layer in GoBOED. Leveraging amortized variation inference, we create a differentiable pipeline that can identify optimal experiments targeting decision value. Unlike traditional information-maximizing designs, GoBOED can provide flexibility in experimental selection, as the experiment with the lowest data acquisition cost may be prioritized when multiple experiments lead to equivalent decision quality despite their difference in reducing the parameter uncertainty. The application of GoBOED to real-world problems, such as epidemic management and pharmacokinetic control, demonstrates the efficacy of our proposed goal-driven experimental design approach.

1 INTRODUCTION

When experiments are expensive, time-consuming, or potentially dangerous, optimizing the experiment becomes crucial. For example, for systems identification or dynamic model learning in such scenarios, we need to carefully select the most informative experiments to accurately estimate the model parameters that govern the system. **Bayesian optimal experimental design (BOED)** provides a systematic framework specifically designed for this purpose, allowing researchers to identify maximally informative experimental designs (Chaloner & Verdinelli, 1995; Rainforth et al., 2024). This approach has found applications across diverse fields including psychology (Bach, 2023), geophysics (Strutz & Curtis, 2024), and other domains where experimental resources are limited.

However, BOED comes with significant computational challenges. It inherently requires simulating numerous scenarios to estimate the posterior distribution of the model parameters, often involving complex calculations such as Kullback-Leibler (KL) divergence between distributions or covariance matrix evaluations. To address these challenges, researchers have proposed various approaches, including nested Monte Carlo methods (Rainforth et al., 2018) and computational frameworks (Ryan et al., 2016). However, these methods require large numbers of samples for convergence, ultimately perpetuating the computational burden. These substantial sampling requirements make BOED computationally expensive in practice, limiting its applicability to complex real-world systems.

To tackle computational bottlenecks while approximating complicated posterior distributions precisely, researchers have developed numerous amortized Bayesian inference methods. These approaches include variational inference (Foster et al., 2019; 2020), normalizing flows (Dong et al., 2025; Orozco et al., 2024), diffusion models (Iollo et al., 2025), and mutual information neural

estimation (Kleinegesse & Gutmann, 2020). In cases where experiments are performed sequentially, reinforcement learning techniques have proven particularly valuable, with several methods proposed to adaptively optimize the experimental design using accumulated information gain for every step (Foster et al., 2021; Ivanova et al., 2021; Blau et al., 2022).

Building on the BOED foundation, recent research has developed computationally efficient strategies for robust decision-making under uncertainty in real-world applications. For example, optimal control algorithms have been shown effective in epidemic management for developing policies that reduce both public health and economic impacts (Linde et al., 2009; Nowzari et al., 2016; Paré et al., 2020; Gardner et al., 2021). By focusing on the structure of the compartmental epidemiological models before evaluating the final output, computationally efficient strategies based on the Susceptible-Infected-Quarantined-Recovered (SIQR) model have been developed for optimal control of infectious disease while maintaining social functioning and minimizing disruptions under resource constraints (Ma et al., 2023; Ofir et al., 2022). Another real-world example is for quantifying how medicines are absorbed, distributed, metabolized, and eliminated in the body using Pharmacokinetic (PK) models (Mould & Upton, 2013; Zou et al., 2020). Within PK modeling, dose optimization is critical to maintain efficacy while minimizing adverse toxicity (Silva et al., 2025). A detailed discussion of related work is provided in Appendix A.

However, all of these models are typically abstracted and significantly simplified, and the parameters in these models are inherently uncertain, making accurate estimation critical for effective decision making. Vitková et al. (2023) addressed parameter uncertainty by analyzing open-loop control cycles. While robust control methods can accommodate parameter uncertainty and more general model uncertainty (Nemirovski, 2012) , they often lead to overly conservative policies. Therefore, reducing uncertainty by BOED is necessary for effective control outcomes.

Considering experimental design for the final operational goal of more effective decision-making under uncertainty, we propose a new **Goal-driven BOED (GoBOED)** framework to strategically design experiments to reduce model uncertainty that most significantly affects decision outcomes. We utilize variational inference methods to approximate accurate posterior distributions (Foster et al., 2019), while employing convex optimization methods for optimal control (Talaei et al., 2024; Ambikapathi et al., 2015). Figure 1 visualizes our proposed GoBOED framework.

Our GoBOED has significantly reduced computational time while maintaining the interpretability of the solution. At the optimal point, we can perform Lagrangian sensitivity analysis and use the derivative information to guide experimental design. This allows us to efficiently identify informative experimental designs that have the greatest impact on decision-making processes, creating a more direct bridge between experimental observation and control implementation in convex settings.

In summary, our main contributions are:

- We propose GoBOED to integrate Bayesian optimal experimental design with robust optimal control under uncertainty governed by convex optimization, enabling effective and efficient robust decision-making under uncertainty.
- We develop computational strategies that efficiently train posterior distributions while simultaneously making decisions under uncertainty.
- The proposed method is applicable to a broader class of real-world control problems formulated as convex optimization problems with uncertainty-aware complex system modeling, with demonstrated performances in epidemic management and pharmacokinetic control.

2 BACKGROUND

2.1 BAYESIAN OPTIMAL EXPERIMENTAL DESIGN (BOED)

BOED is an information-theoretic approach to the problem of identifying which experiments are most informative. It consists of a prior assumption about unknown model parameters $\boldsymbol{\theta} \in \Theta$, design variables $\xi \in \Xi$, a forward model $f : \Theta \times \Xi \to Y$, and observations $y \in Y$ given by $y = f(\boldsymbol{\theta}, \xi) + \epsilon$, where ϵ denotes noise, e.g., $\epsilon \sim \mathcal{N}(0, I)$.

Given these assumptions, we evaluate experimental designs by computing the expected information gain (EIG). This involves calculating the KL divergence between the posterior and prior distribu-

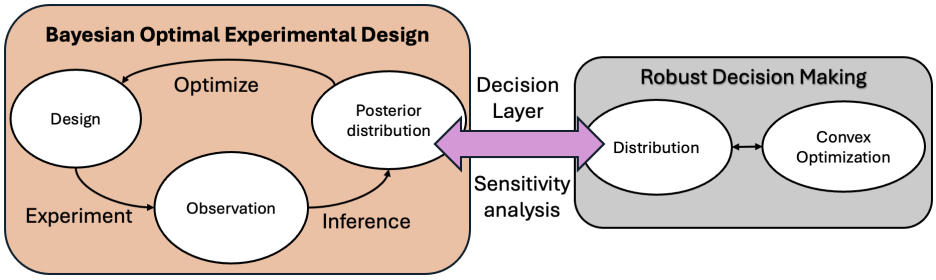

Figure 1: The figure illustrates how Bayesian Optimal Experimental Design (BOED) and robust decision making interact. In BOED (left), observations from an experiment update the posterior, and the (expected) posterior distribution then guides the design of the next experiment. The decision layer (center) maps the posterior to decision-relevant quantities (e.g., losses, risks, constraints) and supports sensitivity analysis of how decisions vary with uncertainty. In the robust decision-making module (right), a convex optimization uses this uncertainty representation to select actions that perform well under uncertainty. Our method couples these components by embedding the robust decision problem within the BOED loop: sensitivity analysis identifies regions of the posterior most relevant for decisions, and the resulting robust decisions are used to prioritize the next experiment.

tions, and taking the expectation of this value with respect to the marginal likelihood. This formulation can be reformulated to estimate EIG by sample average approximation (SAA) using Bayes' rule:

$$EIG(\xi) := \mathbb{E}_{p(\boldsymbol{\theta})p(y|\boldsymbol{\theta},\xi)}\left[\log \frac{p(y|\boldsymbol{\theta},\xi)}{p(y|\xi)}\right] \tag{1}$$

We then find the optimal experimental design $\xi^* = \arg\max_\xi EIG(\xi)$ to identify the most informative experiment.

Since computing the posterior distribution and KL divergence between two distributions can be computationally expensive, EIG has been often efficiently estimated using SAA with approximated posterior distributions. As our paper focuses on real-world scientific problems and decision-aware optimization, we employ variational inference (Foster et al., 2019; 2020). A key benefit of this approach is that EIG with variational inference asymptotically converges to the true value as the number of samples increases, ensuring the soundness of decision outputs alongside computational efficiency.

When using variational inference, we can approximate the EIG by:

$$EIG(\xi) \approx \mathbb{E}_{p(\boldsymbol{\theta})p(y|\boldsymbol{\theta},\xi)}\left[\log p(y|\boldsymbol{\theta},\xi) - \mathbb{E}_{q_\phi(\boldsymbol{\theta}|y,\xi)}\left[\log \frac{p(y|\boldsymbol{\theta},\xi)q_\phi(\boldsymbol{\theta}|y,\xi)}{p(\boldsymbol{\theta})}\right]\right], \tag{2}$$

where $q_\phi(\boldsymbol{\theta}|y,\xi)$ is an approximated posterior distribution, and ϕ represents neural network parameters for the variational inference network, which generates the posterior parameters based on observation and design.

2.2 From BOED to Robust Decision-Making Under Model Uncertainty

In many real-world applications, optimal decisions depend on uncertain model parameters $\boldsymbol{\theta}$. Standard robust decision-making accounts for this uncertainty when choosing an action, but here we emphasize a complementary lever: *BOED* to actively reduce the uncertainty that matters for downstream decisions.

We first formalize the robust decision problem given a posterior. After running a design ξ and observing y, we obtain a posterior distribution over parameters $\boldsymbol{\theta}$, $p(\boldsymbol{\theta} \mid y, \xi)$. Let $J(\boldsymbol{q};\boldsymbol{\theta})$ denote the application cost corresponding to decision-making $\boldsymbol{q}$ given inferred model parameters $\boldsymbol{\theta}$ (equivalently, $J(\boldsymbol{q} \mid y,\xi) = \mathbb{E}_{\boldsymbol{\theta}\sim p(\boldsymbol{\theta}|y,\xi)}[\,J(\boldsymbol{q};\boldsymbol{\theta})\,]$). We seek a decision that is robust with respect to the updated posterior:

$$\boldsymbol{q}^\star \in \arg\min_{\boldsymbol{q}\in\mathcal{Q}} \ \rho_{\boldsymbol{\theta}\sim p(\boldsymbol{\theta}|y,\xi)}\big[J(\boldsymbol{q};\boldsymbol{\theta})\big], \tag{3}$$

where ρ is a risk functional (e.g., expectation or chance constraints).

BOED chooses the design before observing y so that the posterior—hence the robust action in equation 3—improves in the directions that matter for $J(\boldsymbol{\theta}; \boldsymbol{q})$. The combining objective is

$$\xi^\star \in \arg\min_{\xi \in \Xi} \; \mathbb{E}_{p(y|\xi)}\Big[\min_{\boldsymbol{q} \in \mathcal{Q}} \; \rho_{\boldsymbol{\theta} \sim p(\boldsymbol{\theta}|y,\xi)}\big[J(\boldsymbol{q}; \boldsymbol{\theta})\big]\Big], \tag{4}$$

which is decision-focused: it prioritizes experiments that most reduce eventual robust loss, not just overall parameter uncertainty. We illustrate with two applications: epidemic management and pharmacokinetic control.

Optimal control for epidemiology Following Talaei et al. (2024), we study robust epidemiological management based on a compartmental SIQR model. The population is divided into susceptible (s), asymptomatic infected (x^a), symptomatic infected (x^s), and recovered (h). Let $\boldsymbol{d} = [\, s,\, x^a,\, x^s,\, h\,]^\top$ denote the state vector. The dynamics are

$$\dot{\boldsymbol{d}} = M(\boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{q})\, \boldsymbol{d},$$

where M represents the system matrix modeling epidemic infection dynamics, $\boldsymbol{\beta} = (\beta^a, \beta^s)$ are the transmission rates for asymptomatic and symptomatic cases, $\boldsymbol{\gamma} = (\gamma^a, \gamma^s)$ are the corresponding recovery rates, and $\boldsymbol{q} = (q^a, q^s)$ are the quarantine rates. Here $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$ are uncertain.

We optimize the quarantine strategy as described in Talaei et al. (2024). The objective function that minimizes economic costs:

$$\min_{\boldsymbol{q} \in [0,1)^2} \left\{ J(\boldsymbol{q}; \boldsymbol{\beta}, \boldsymbol{\gamma}) = \left(\frac{z^a}{1 - q^a} + \frac{z^s}{1 - q^s} \right) \right\} \quad \text{s.t. } \lambda_{\max}(M(\boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{q})) \leq -\alpha,$$

where z^a represents the economic cost for asymptomatic quarantine, z^s represents the economic cost for symptomatic quarantine, $\lambda_{\max}$ represents the largest eigenvalue of the matrix, and $\alpha > 0$ is a constraint ensuring the stability of the system. We adopt the quarantine cost function proposed by Talaei et al. (2024). Note that the objective function of this convex programming formulation $J(\boldsymbol{q}; \boldsymbol{\beta}, \boldsymbol{\gamma})$ is not directly dependent on SIQR model parameters. The detailed derivation of solution can be found in Appendix D.1

Optimal dosing with pharmacokinetic models Considering pharmacokinetic (PK) control as another example, the concentration at time t for a drug administered orally can be modeled using the Bateman function (Bateman, 1910):

$$y(t) = \frac{D \cdot q^c \cdot k_a}{V \cdot (k_a - k_e)} \left(e^{-k_e \cdot t} - e^{-k_a \cdot t}\right) (1 + \epsilon_{\text{mult}}) + \epsilon_{\text{add}},$$

where V is the volume of distribution, q^c is the dosing rate, k_a is the absorption rate constant, k_e is the elimination rate constant, D is the dose administered, ϵ_{mult} is a multiplicative error term, and ϵ_{add} is an additive error term. This formulation captures the dynamics of drug absorption and elimination.

Given the drug's potential toxicity, dosing should maintain systemic exposure within the therapeutic window—avoiding toxic concentrations while not falling below the minimum effective concentration. We consider the maximum concentration $C_{\max}$ and the area under the concentration curve (AUC) as constraints and define a convex cost function $J(q)$ to discourage high dosing, where C_1 and C_2 are constant. The constrained problem can be written in the following forms,

$$\min_{q^c \in [0,1]} \quad C_1 \cdot q^c + C_2 \cdot (q^c)^2 \quad \text{s.t.} \quad y_{\max}(q^c; \boldsymbol{\theta}) \leq y_{\text{thresh}}, \quad \text{AUC}(q^c; \boldsymbol{\theta}) \;\geq\; \text{AUC}_{\min}.$$

Here $\boldsymbol{\theta}$ denotes $[k_a, k_e, V]$ in the Bateman function, and the solution to this optimization formulation can be found in Appendix D.2. We formulate an analogous convex optimization problem for the PK model, mirroring the SIQR formulation.

3 METHODS

We develop an integrated framework, GoBOED, for goal-driven experimental design by bringing together BOED (introduced in eq. (2)) and robust optimal control under model uncertainty through convex optimization (described in eq. (3)). Our primary objective is to identify an optimal experimental design ξ^* that minimizes the expected controlled economic cost (i.e., the control objective),

where the expectation is taken over posterior distributions of the parameters $\boldsymbol{\theta}$. This approach allows us to update our beliefs with new observations and improve decision-making regarding goal-driven optimal experiment and robust optimal control based on an uncertain model, following the methodology of Chaloner & Verdinelli (1995).

For a given experimental design ξ and observed data y, we formulate the optimization problem as:

$$\min_{\boldsymbol{q}} \ \mathbb{E}_{p(\boldsymbol{\theta}|y,\xi)}[J(\boldsymbol{q}) \quad \text{s.t. Constraints}(\boldsymbol{\theta})] \,. \tag{5}$$

Here, $J(\boldsymbol{q})$ is the control cost, assumed convex in $\boldsymbol{q}$, and it does not depend on $\boldsymbol{\theta}$ directly; uncertainty enters through the constraints in eq. (5). Under this convexity assumption we can carry out sensitivity analysis of the optimizer with respect to $\boldsymbol{q}$. The uncertain model parameters are collected in $\boldsymbol{\theta} \sim p(\boldsymbol{\theta} \mid y, \xi)$, and the key challenge is evaluating the chance term from the posterior. We address this using importance sampling with a variational proposal as detailed below.

To approximate the posterior distribution, we employ stochastic variational inference. Specifically, we approximate the true posterior $p(\boldsymbol{\theta}|y,\xi)$ using a variational distribution $q_\phi(\boldsymbol{\theta}|y,\xi)$, where ϕ denotes the variational parameters. The variational distribution is optimized by maximizing the evidence lower bound (ELBO), which can be expressed as:

$$\mathcal{L}_{VI}(\phi; y) = \mathbb{E}_{q_\phi(\boldsymbol{\theta}|y,\xi)} \left[\log p(y|\boldsymbol{\theta},\xi)\right] - D_{KL}\left(q_\phi(\boldsymbol{\theta}|y,\xi) \| p(\boldsymbol{\theta})\right). \tag{6}$$

Once we have obtained the variational posterior $q_\phi(\boldsymbol{\theta}|y,\xi)$, we use it to estimate the expectation of the constraints. Since q_ϕ may not exactly match the true posterior, we apply importance sampling to correct for the discrepancy. For any function $f(\boldsymbol{\theta})$, the expectation under the true posterior can be estimated as

$$\mathbb{E}_{p(\boldsymbol{\theta}|y,\xi)}[f(\boldsymbol{\theta})] \approx \frac{\sum_{i=1}^N f(\boldsymbol{\theta}_i) w(\boldsymbol{\theta}_i)}{\sum_{i=1}^N w(\boldsymbol{\theta}_i)}, \tag{7}$$

where $(\boldsymbol{\theta}_i)$ are samples from $q_\phi(\boldsymbol{\theta}|y,\xi)$, and $w(\boldsymbol{\theta}_i) = \frac{p(y|\boldsymbol{\theta}_i,\xi)p(\boldsymbol{\theta}_i)}{q_\phi(\boldsymbol{\theta}_i|y,\xi)}$ are importance weights.

In particular, for the constraint, we compute Constraints$(\mathbb{E}_{p(\boldsymbol{\theta}|y,\xi)}\boldsymbol{\theta})]$ using the above formula with $f(\boldsymbol{\theta}) = \boldsymbol{\theta}$. This allows us to evaluate the constraint and solve the optimization problem for each ξ and y. Thus, our framework integrates BOED with optimal control under updated posterior distribution uncertainty, leveraging variational inference and importance sampling.

3.1 Alternative formulation using chance constraints

To better incorporate the uncertainty in parameters $\boldsymbol{\theta}$ for robust robust optimization, we apply an alternative formulation using chance constraints, a powerful tool in optimization under uncertainty ensuring that critical conditions hold with a specified probability (Charnes & Cooper, 1959; Miller & Wagner, 1965). Given that the objective function is independent of $\boldsymbol{\theta}$ in our formulations, we can concentrate on the constraints with posterior samples $\boldsymbol{\theta}_i \sim p(\boldsymbol{\theta}|y,\xi), i = 1, ..., N$. Directly imposing constraints for each sample would be overly restrictive. Instead, we introduce a chance constraint to ensure that each sample from the updated posterior distribution satisfies a stability condition with high probability.

The alternative formulation, denoted $\mathcal{L}_{\text{OC-CC}}$, is defined as:

$$\min_{\boldsymbol{q}} \ J(\boldsymbol{q}) \quad \text{s.t. } \mathbb{P}(\text{Constraints}(\boldsymbol{\theta}) \mid y, \xi) \geq \eta \tag{8}$$

where $\boldsymbol{\theta}$ comes from the posterior distribution $p(\boldsymbol{\theta}|y,\xi)$, and $\eta \in (0,1)$ is the desired probability level.

To evaluate the chance term in eq. (8) efficiently, we use importance sampling with a variational proposal. Specifically, we set $f(\boldsymbol{\theta}) = \mathbf{1}\{\text{Constraints}(\boldsymbol{\theta})|y,\xi\}$ in eq. (7) to estimate $\mathbb{P}\big(\text{Constraints}(\boldsymbol{\theta}) \mid y, \xi\big)$ stably.

1. Draw N samples $\{\boldsymbol{\theta}_i\}_{i=1}^N$ from $q_\phi(\boldsymbol{\theta} \mid y, \xi)$.
2. Compute importance weights, and normalize it

$$w_i \ = \ \frac{p(\boldsymbol{\theta}_i)\, p(y \mid \boldsymbol{\theta}_i, \xi)}{q_\phi(\boldsymbol{\theta}_i \mid y, \xi)}, \qquad \tilde{w}_i \ = \ \frac{w_i}{\sum_{j=1}^N w_j}.$$

3. Estimate the posterior probability (given y, ξ):

$$\mathbb{P}(\text{Constraints}(\boldsymbol{\theta}) \mid y, \xi) \;=\; \sum_{i=1}^{N} \tilde{w}_i \, \mathbf{1}\{\text{Constraints}(\boldsymbol{\theta}) \mid y, \xi\}\,.$$

Using these importance-weighted samples, we impose the constraints and solve the resulting optimization via convex optimization.

3.2 Robust control using Conditional Value-at-Risk (CVaR)

As an alternative to the scenario-based chance constraint in Appendix E.1.1, we control the *tail* of constraint violations using Conditional Value-at-Risk (CVaR).

Let the feasibility event be $\mathbb{P}(\text{Constraints}(\boldsymbol{\theta}) \mid y, \xi) \geq \eta$. For samples $\{\boldsymbol{\theta}_i\}_{i=1}^{N}$ from the posterior distribution, define the per-sample violation $v_i(\boldsymbol{q}; \boldsymbol{\theta}_i)$. We enforce

$$\mathrm{CVaR}_{\eta}\big(v(\boldsymbol{q}; \boldsymbol{\theta})\big) \leq 0 \tag{9}$$

using the Rockafellar–Uryasev sample-average form with optional normalized weights $\bar{w}_i$:

$$s_i \;\geq\; v_i(\boldsymbol{q}; \boldsymbol{\theta}_i) - \tau, \quad i = 1, \ldots, N, \qquad \tau + \frac{1}{1-\eta} \sum_{i=1}^{N} \bar{w}_i \, s_i \;\leq\; 0, \tag{10}$$

with decision variables $\tau \in \mathbb{R}$ and $s_i \geq 0$. At optimality, $s_i = (v_i(\boldsymbol{q}; \boldsymbol{\theta}_i) - \tau)_+$, so only the upper tail beyond the η-quantile contributes.

3.3 Differentiable Decision layer

We embed the robust decision problem from Appendix B as a *differentiable decision layer*. Given parameters $\boldsymbol{\theta}$, the layer solves the convex program and returns the optimal pair $(\hat{\boldsymbol{q}},\, J^*(\hat{\boldsymbol{q}}; \boldsymbol{\theta}))$ together with the KKT multipliers $\hat{\boldsymbol{\lambda}}$ of the active constraints. Under standard regularity, $J^*(\hat{\boldsymbol{q}}; \boldsymbol{\theta})$ is differentiable, and KKT sensitivity yields

$$\nabla_{\boldsymbol{\theta}} J^*(\boldsymbol{\theta}) = \partial_{\boldsymbol{\theta}} J(\hat{\boldsymbol{q}}; \boldsymbol{\theta}) \;+\; \sum_{i \in \mathcal{A}} \hat{\lambda}_i \, \nabla_{\boldsymbol{\theta}} \, g_i(\hat{\boldsymbol{q}}, \boldsymbol{\theta}; y, \xi),$$

where $g_i(\boldsymbol{q}, \boldsymbol{\theta}; \xi) \leq 0$ are the constraints and $\mathcal{A}$ is the active set. (If J does not depend on $\boldsymbol{\theta}$, the first term vanishes.)

With this decision layer, gradients with respect to the variational parameters ϕ and the design ξ follow by the chain rule:

$$\frac{dJ^*}{d\psi} = \big(\nabla_{\boldsymbol{\theta}} J^*(\boldsymbol{\theta})\big)^{\top} \frac{\partial \boldsymbol{\theta}}{\partial \psi}, \qquad \psi \in \{\phi, \xi\},$$

with the reparameterization $\boldsymbol{\theta} = h(\epsilon; y, \xi, \phi)$, $\epsilon \sim p(\epsilon)$ (see Section 3.5 for details). Any explicit penalties that depend on ξ through the constraints are handled via $\nabla_{\xi} h_i(\hat{\boldsymbol{q}}, \boldsymbol{\theta}; y, \xi)$.

In summary, the decision layer solves the inner problem once (forward) and differentiates via KKT (backward), making the robust decision step directly compatible with end-to-end training. For implementation, we use `cvxpylayers` Agrawal et al. (2019) to map the convex optimization result directly to the target output.

3.4 Formulation of the optimization problem

By combining convex optimization with stochastic variational inference in the previous section, we achieve robust decision making under model parameter uncertainty. Our next objective is to determine the optimal experimental design ξ^* and corresponding variational parameters ϕ^* that jointly minimize the expected cost of decision making while ensuring accurate approximation of the posterior distribution for parameters $\boldsymbol{\theta}$. We compute the expectation over the marginal likelihood,

following approaches similar to Krishnan & Tickoo (2020a); Lacoste–Julien et al. (2011), because the posterior distribution depends on a specific observation.

We formulate our optimization problem as:

$$(\xi^*, \phi^*) = \arg\min_{\xi \in \Xi, \phi} \mathbb{E}_{p(y|\xi)} \left[\mathcal{L}_{OC}(y, \xi; \phi) - \mathcal{L}_{\text{VI}}(\phi; y, \xi) \right], \tag{11}$$

where $\mathcal{L}_{OC} = \min_{\boldsymbol{q}} \mathbb{E}_{q_\phi(\boldsymbol{\theta}|y,\xi)}[J(\boldsymbol{q}) + \text{Constraints}(\boldsymbol{\theta})]$.

3.4.1 Gradient estimation

To optimize ϕ, we use reparameterization (Burda et al., 2015; Foster et al., 2020) to stabilize training and to differentiate the variational objective w.r.t. ϕ. We maximize eq. (6) by writing $\boldsymbol{\theta}_i = g(\epsilon_i; y, \xi, \phi)$, where g represents variational encoder, and $\epsilon_i \sim p(\epsilon)$ is a standard normal random variable. The gradient is approximated via Monte Carlo estimation as described in (Foster et al., 2019)

For design variable ξ, we set our robust decision–making loss function $L(\xi) = \mathbb{E}_{p(y|\xi)}[\mathcal{L}_{OC}(y, \xi; \phi)]$. By chain rule, the gradient with respect to ξ can be written in the following forms,

$$\frac{\partial L}{\partial \xi} = \mathbb{E}_{p(y|\xi)} \left[\frac{\partial \mathcal{L}_{OC}(y, \xi; \phi)}{\partial \xi} + \mathcal{L}_{OC}(y, \xi; \phi) \frac{\partial \log p(y|\xi)}{\partial \xi} \right] \tag{12}$$

The first term, $\frac{\partial \mathcal{L}_{OC}(y,\xi;\phi)}{\partial \xi}$, can be computed using implicit differentiation through decision layer, as described in Section 3.3.

The second term, $\mathcal{L}_{OC}(y, \xi; \phi) \frac{\partial \log p(y|\xi)}{\partial \xi}$, is straightforward to compute once an observation model (e.g., Poisson or Gaussian noise) is specified.

3.5 Optimization of experimental design

When the forward model is expensive to solve, recomputing a variational posterior at every candidate design is prohibitive. We therefore train a single amortized variational network that maps a design–observation pair (ξ, y) to the parameters of a posterior $q_\phi(\boldsymbol{\theta} \mid \xi, y)$ in one shot, and then use this network for gradient-based design optimization. This enables highly efficient experimental design.

Amortized VI (one-shot training). The amortizer takes (ξ, y) as input, lifts them into a shared latent space of width d, and fuses the signals with a single-head attention block (queries from y, keys from ξ, values from both). $q_\phi(\boldsymbol{\theta} \mid \xi, y)$ (see Appendix C for architecture). We train ϕ once using simulated trajectories $\{\xi_t, y_t\}_{t \in \mathcal{T}}$ and maximize the summed ELBO,

$$\hat{\mathcal{L}}_{\text{VI}}(\phi) \;=\; \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \mathbb{E}_{\boldsymbol{\theta} \sim q_\phi(\cdot|\xi_t, y_t)} \Big[\log p(y_t \mid \boldsymbol{\theta}, \xi_t) + \log p(\boldsymbol{\theta}) - \log q_\phi(\boldsymbol{\theta} \mid \xi_t, y_t) \Big],$$

which corresponds to eq. (6) evaluated across all $t \in \mathcal{T}$. The attention layer aggregates information across multiple designs without discarding earlier contributions, avoiding reliance on a fixed handcrafted summary.

Design gradients via reparameterization. After training, the amortizer produces variational parameters $\gamma_\phi(\xi, y)$ for a family $q_\phi(\boldsymbol{\theta} \mid \xi, y)$. Assuming the family is reparameterizable, samples can be written as a deterministic transformation of base noise:

$$\boldsymbol{\theta}_s(\xi, y) \;=\; h_\phi\big(\varepsilon_s;\, \gamma_\phi(\xi, y),\, \xi, y\big), \qquad \varepsilon_s \sim r(\varepsilon) \text{ (e.g., } \mathcal{N}(0, I)).$$

This yields pathwise derivatives $\partial \boldsymbol{\theta}_s(\xi, y)/\partial \xi$, which we propagate through the decision layer (see Section 3.3).

Computational benefit. This procedure concentrates simulation cost in a single offline training phase. Once q_ϕ is learned, design optimization only involves forward passes through the amortizer and standard autodiff, yielding efficient and accurate gradients for ξ without repeatedly solving the forward model or re-fitting a variational posterior.

4 RESULTS

We present numerical experiments on robust decision-making for goal-driven experimental design for two use cases, in order to demonstrate the effiacy and the general applicability of the proposed GoBOED framework: *epidemic management (SIQR)* and *pharmacokinetic (PK) control*, both under model uncertainty. Detailed model parameters and solver settings are provided in Appendix F.

To compare GoBOED with standard BOED baselines, we study the problem of choosing a single observation time. Let T denote the time horizon and let $\boldsymbol{\xi} \in \{0,1\}^T$ be a one-hot design vector with $\sum_{t=1}^{T} \xi_t = 1$; the unique index $t^\star$ with $\xi_{t^\star} = 1$ is the chosen measurement time. We collect exactly one measurement at $t^\star$ and use this datum to update the posterior over model parameters. Because both the SIQR and PK settings are time-indexed, this formulation applies to both. At $t^\star$, we observe the counts of asymptomatic and symptomatic infections for the SIQR model and the drug concentration in blood for the PK model; measurement noise is modeled as Poisson or Gaussian, respectively, depending on the data modality.

Our goal is to select the single measurement time that best supports downstream decision-making. Traditional BOED via EIG targets maximal reduction in parameter uncertainty, whereas the robust optimal control objective is sensitive to particular parameter combinations and system dynamics. Accordingly, for each candidate design $\boldsymbol{\xi}$ we compute and visualize both the EIG and the robust optimal control cost over the entire design space. The resulting design objective surfaces for the SIQR (top) and PK models (bottom) are shown in Figure 2.

We estimate EIG (cf. eq. (1)) using nested Monte Carlo with $5{,}000$ outer samples (over y) and $3{,}000$ inner samples for the marginal likelihood. The BOED-selected optimal observation times are day 5 for the SIQR model and hour 17 for the PK model.

Robust optimal control in the presence of model uncertainty. We solve the chance-constrained problem in eq. (8), enforcing a 90% probability of constraint satisfaction under the posterior induced by a given observation time. Empirically, designs with larger EIG reduce constraint uncertainty and thereby yield lower optimal cost in both examples. For SIQR, the objective is relatively flat for observation times between days 4 and 8; for PK, observing later (roughly 15–23 hours) reduces the dose required to meet the therapeutic targets. Unless otherwise noted, we draw 500 datasets y and, for each y, 40 posterior samples of $\boldsymbol{\theta}$ (128 for PK).

CVaR-based constraints. We also consider the CVaR formulation in eq. (9): for SIQR we constrain the CVaR of the dominant eigenvalue, and for PK we constrain the CVaR of $C_{\max}$ relative to its threshold. We set the level to $\alpha = 0.9$ (controlling the expected violation over the worst 10% of posterior realizations). The qualitative trends mirror those under chance constraints: higher-EIG designs generally achieve lower robust cost; SIQR exhibits a plateau around days 4–8, and in PK a later observation (e.g., near 24 hours) further reduces the dose required to achieve the target performance. We use the same sampling budgets as above (500 draws of y; 40 posterior draws per y, or 128 for PK).

These results highlight a key insight: maximizing EIG alone does not necessarily optimize constraint-aware economic performance. Instead, our approach identifies a broad near-optimal window for SIQR—approximately days 4–8—within which observation timing can be chosen with minimal loss in both information gain and economic efficiency. For PK dose optimization, we similarly find a wide near-optimal window (about 15–23 hours). Within this window, latter observations (toward 22–23 hours) generally yield robust control with lower cost than the BOED-optimal 17 hours while achieving comparable EIG. This scheduling flexibility is practically valuable when exact timing for observation measurements is restricted due to logistical or operational constraints.

Gradient-based search over discrete times. We treat the observation time as a scalar $t \in \{1, \dots, T\}$ (equivalently, a one-hot $\boldsymbol{\xi}(t)$) and directly optimize t using automatic differentiation (AD). Starting from a mid-point integer $t_0 = \lfloor (T+1)/2 \rfloor$, we compute the AD gradient of the robust objective with respect to t and take projected gradient steps:

$$\ell_k \;=\; \frac{\partial J(t)}{\partial t}\Big|_{t=t_k}, \qquad t_{k+1} \;=\; \Pi_{[1,T]}\big(t_k - \eta_k\, \ell_k\big), \qquad \hat{t}_{k+1} \;=\; \mathrm{round}(t_{k+1}),$$

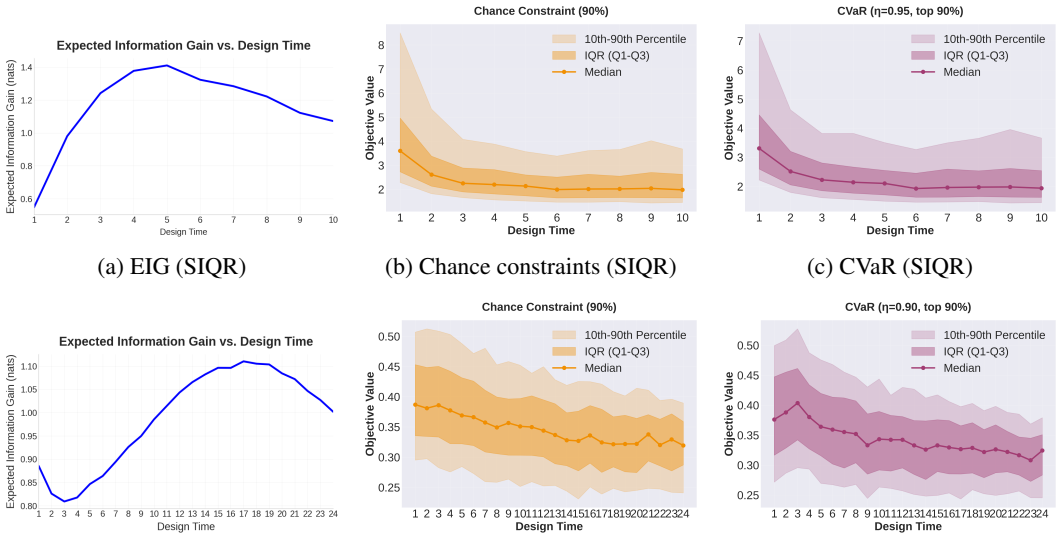

(a) EIG (SIQR) (b) Chance constraints (SIQR) (c) CVaR (SIQR)

(d) EIG (PK) (e) Chance constraints (PK) (f) CVaR (PK)

Figure 2: Comparison of experimental design metrics and control strategies across two models. Top row: SIQR epidemiological model—(a) expected information gain (EIG) over observation time ξ; (b) expected optimal cost under chance constraints with confidence level $\eta = 0.9$; (c) expected optimal cost under CVaR. Bottom row: pharmacokinetic (PK) model—(d) EIG; (e) chance constraints ($\eta = 0.9$); (f) CVaR. The horizontal axis is observation time ξ; the vertical axis shows EIG or expected optimal control costs. While the BOED-optimal design typically pinpoints a specific time, the goal-driven robust objective admits a broader near-optimal window—offering greater scheduling flexibility under real-world constraints.

where $\Pi_{[1,T]}$ clips to the feasible interval and η_k is the step size. After each update, we evaluate the design at the integer index $\hat{t}_{k+1}$ by recomputing the posterior induced by a single observation at $\hat{t}_{k+1}$ and then re-evaluating the robust optimal cost $J(\hat{t}_{k+1})$. We repeat these steps until convergence.

For the SIQR model, the proposed procedure converges near the BOED design by day 5, with days 6–7 showing similarly small gradient norms. For the PK model, it selects the 22-hour design for both the CVaR and chance-constrained criteria, which aligned with the plot Figure 2. We also confirmed a monotonic decrease in the gradient norm, consistent with the trend observed in the associated plot.

5 Conclusion

We developed a computational methodology that jointly trains a posterior approximation and optimizes uncertainty-aware, goal-driven experimental decisions that enables robust optimal management/control under model uncertainty. By bridging convex optimal control with BOED, the framework supports real-world deployments that could meaningfully influence epidemic response strategies and dosing optimization while improving computational efficiency. Applying the method to large-scale epidemiology, clinical, and drug discovery datasets—and integrating it into production workflows–remains important future work.

Limitations. Our framework assumes convex optimization, which enables differentiation through the control layer via Lagrangian sensitivity analysis. This assumption is both a strength and a constraint: for non-convex objectives we typically obtain only locally optimal solutions, increasing optimization difficulty. Performance also depends on the quality of variational inference; inaccuracies in the learned posterior can degrade decisions and undermine reliability, particularly in real-world settings. Future work should strengthen robustness and expand posterior expressivity—for example, by leveraging generative families such as diffusion and flow based models.

LLM Usage: This manuscript was copy-edited for grammar and style using ChatGPT (OpenAI; accessed September 2025). The authors drafted all text and iteratively reviewed and revised AI-

suggested edits. All ideas, methods, analyses, and conclusions were developed solely by the authors, who accept full responsibility for the final content. No confidential or reviewer-only material was provided to the model.

REFERENCES

Bram C Agema, Birgit CP Koch, Ron HJ Mathijssen, and Stijn LW Koolen. From prospective evaluation to practice: model-informed dose optimization in oncology. *Drugs*, 85(4):487–503, 2025.

Akshay Agrawal, Brandon Amos, Shane Barratt, Stephen Boyd, Steven Diamond, and J Zico Kolter. Differentiable convex optimization layers. *Advances in neural information processing systems*, 32, 2019.

ArulMurugan Ambikapathi, Tsung-Han Chan, Chia-Hsiang Lin, Fei-Shih Yang, Chong-Yung Chi, and Yue Wang. Convex-optimization-based compartmental pharmacokinetic analysis for prostate tumor characterization using dce-mri. *IEEE Transactions on Biomedical Engineering*, 63(4): 707–720, 2015.

MOSEK ApS. *The MOSEK Python Fusion API manual. Version 11.0.*, 2025. URL `https://docs.mosek.com/latest/pythonfusion/index.html`.

Ahmed Attia, Alen Alexanderian, and Arvind K Saibaba. Goal-oriented optimal design of experiments for large-scale Bayesian linear inverse problems. *Inverse Problems*, 34(9):095009, 2018.

Dominik R Bach. Experiment-based calibration in psychology: Optimal design considerations. *Journal of Mathematical Psychology*, 117:102818, 2023.

Harry Bateman. The solution of a system of differential equations occurring in the theory of radioactive transformations. In *Proc. Cambridge Philos. Soc.*, volume 15, pp. 423–427, 1910.

Freddie Bickford Smith, Andreas Kirsch, Sebastian Farquhar, Yarin Gal, Adam Foster, and Tom Rainforth. Prediction-oriented Bayesian active learning. In Francisco Ruiz, Jennifer Dy, and Jan-Willem van de Meent (eds.), *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, volume 206 of *Proceedings of Machine Learning Research*, pp. 7331–7348. PMLR, 25–27 Apr 2023. URL `https://proceedings.mlr.press/v206/bickfordsmith23a.html`.

Tom Blau, Edwin V. Bonilla, Iadine Chades, and Amir Dezfouli. Optimizing sequential experimental design with deep reinforcement learning. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (eds.), *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 2107–2128. PMLR, 17–23 Jul 2022. URL `https://proceedings.mlr.press/v162/blau22a.html`.

Wolfgang Bock and Yashika Jayathunga. Optimal control and basic reproduction numbers for a compartmental spatial multipatch dengue model. *Mathematical Methods in the Applied Sciences*, 41(9):3231–3245, 2018.

Yuri Burda, Roger Grosse, and Ruslan Salakhutdinov. Importance weighted autoencoders. *arXiv preprint arXiv:1509.00519*, 2015.

Kathryn Chaloner and Isabella Verdinelli. Bayesian experimental design: A review. *Statistical Science*, 10(3):273–304, 1995. doi: 10.1214/ss/1177009939.

Abraham Charnes and William W Cooper. Chance-constrained programming. *Management science*, 6(1):73–79, 1959.

Jiayuan Dong, Christian Jacobsen, Mehdi Khalloufi, Maryam Akram, Wanjiao Liu, Karthik Duraisamy, and Xun Huan. Variational Bayesian optimal experimental design with normalizing flows. *Computer Methods in Applied Mechanics and Engineering*, 433:117457, 2025.

Chinwendu Enyioha, Ali Jadbabaie, Victor Preciado, and George J Pappas. Distributed resource allocation for epidemic control. *arXiv preprint arXiv:1501.01701*, 2015.

Adam Foster, Martin Jankowiak, Eli Bingham, Paul Horsfall, Yee Whye Teh, Tom Rainforth, and Noah Goodman. Variational Bayesian optimal experimental design. *arXiv preprint arXiv:1903.05480*, 2019.

Adam Foster, Martin Jankowiak, Matthew O'Meara, Yee Whye Teh, and Tom Rainforth. A unified stochastic gradient approach to designing Bayesian-optimal experiments. In *International Conference on Artificial Intelligence and Statistics*, pp. 2959–2969. PMLR, 2020.

Adam Foster, Desi R Ivanova, Ilyas Malik, and Tom Rainforth. Deep adaptive design: Amortizing sequential Bayesian experimental design. In *International conference on machine learning*, pp. 3384–3395. PMLR, 2021.

Lauren Gardner, Ensheng Dong, and Hongru Du. Covid-19 dashboard. Technical report, Johns Hopkins University, 2021.

Jonathan Gordon, John Bronskill, Matthias Bauer, Sebastian Nowozin, and Richard E Turner. Meta-learning probabilistic inference for prediction. *arXiv preprint arXiv:1805.09921*, 2018.

Mikhail Hayhoe, Francisco Barreras, and Victor M. Preciado. Multitask learning and nonlinear optimal control of the covid-19 outbreak: A geometric programming approach. *Annual Reviews in Control*, 52:495–507, 2021.

Ashish R. Hota, Jayash Godbole, and Philip E. Paré. A closed-loop framework for inference, prediction, and control of SIR epidemics on networks. *IEEE Transactions on Network Science and Engineering*, 8(3):2262–2278, 2021.

Daolang Huang, Yujia Guo, Luigi Acerbi, and Samuel Kaski. Amortized Bayesian experimental design for decision-making. *Advances in Neural Information Processing Systems*, 37:109460–109486, 2024.

Jacopo Iollo, Christophe Heinkelé, Pierre Alliez, and Florence Forbes. Bayesian experimental design via contrastive diffusions. In *The Thirteenth International Conference on Learning Representations*, 2025.

Desi R Ivanova, Adam Foster, Steven Kleinegesse, Michael U Gutmann, and Thomas Rainforth. Implicit deep adaptive design: Policy-based experimental design without likelihoods. *Advances in Neural Information Processing Systems*, 34:25785–25798, 2021.

Ali Khanafer and Tamer Başar. An optimal control problem over infected networks. In *Proceedings of the International Conference of Control, Dynamic Systems, and Robotics*, volume 125, pp. 1–6, 2014.

Steven Kleinegesse and Michael U Gutmann. Bayesian experimental design for implicit models by mutual information neural estimation. In *International Conference on Machine Learning*, pp. 5316–5326. PMLR, 2020.

Steven Kleinegesse and Michael U Gutmann. Gradient-based Bayesian experimental design for implicit models using mutual information lower bounds. *arXiv preprint arXiv:2105.04379*, 2021.

Ranganath Krishnan and Omesh Tickoo. Improving model calibration with accuracy versus uncertainty optimization. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 18237–18248. Curran Associates, Inc., 2020a. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/d3d9446802a44259755d38e6d163e820-Paper.pdf.

Ranganath Krishnan and Omesh Tickoo. Improving model calibration with accuracy versus uncertainty optimization. *Advances in Neural Information Processing Systems*, 33:18237–18248, 2020b.

Simon Lacoste–Julien, Ferenc Huszár, and Zoubin Ghahramani. Approximate inference for the loss-calibrated Bayesian. In Geoffrey Gordon, David Dunson, and Miroslav Dudík (eds.), *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, volume 15 of *Proceedings of Machine Learning Research*, pp. 416–424, Fort Lauderdale, FL, USA, 11–13 Apr 2011. PMLR. URL https://proceedings.mlr.press/v15/lacoste_julien11a.html.

Yurong Lai, Xiaoyan Chu, Li Di, Wei Gao, Yingying Guo, Xingrong Liu, Chuang Lu, Jialin Mao, Hong Shen, Huaping Tang, et al. Recent advances in the translation of drug metabolism and pharmacokinetics science for drug discovery and development. *Acta Pharmaceutica Sinica B*, 12 (6):2751–2777, 2022.

Sunmi Lee, Gerardo Chowell, and Carlos Castillo-Chávez. Optimal control for pandemic influenza: The role of limited antiviral treatment and isolation. *Journal of Theoretical Biology*, 265(2): 136–150, 2010.

Annika Linde, Maria Rotzén-Ostlund, Benita Zweygberg-Wirgart, Slavica Rubinova, and Mia Brytting. Does viral interference affect spread of influenza? *Eurosurveillance*, 14(40):19354, 2009.

Fei Liu and Michael Buss. Optimal control for heterogeneous node-based information epidemics over social networks. *IEEE Transactions on Control of Network Systems*, 7(3):1115–1126, 2020.

Qianqian Ma, Yang-Yu Liu, and Alex Olshevsky. Optimal fixed lockdown for pandemic control. *IEEE Transactions on Automatic Control*, 69(7):4538–4553, 2023.

Madhusudan Madhavan, Alen Alexanderian, Arvind K Saibaba, Bart van Bloemen Waanders, and Rebekah D White. A control-oriented approach to optimal sensor placement. *arXiv preprint arXiv:2502.15062*, 2025.

Vinh S. Mai, Abdella Battou, and Kevin Mills. Distributed algorithm for suppressing epidemic spread in networks. *IEEE Control Systems Letters*, 2(3):555–560, 2018.

Piet Van Mieghem, Dragan Stevanović, Fernando Kuipers, Cong Li, Ruud Van De Bovenkamp, Daijie Liu, and Huijuan Wang. Decreasing the spectral radius of a graph by link removals. *Physical Review E*, 84(1):016101, 2011.

Bruce L Miller and Harvey M Wagner. Chance constrained programming with joint constraints. *Operations research*, 13(6):930–945, 1965.

Diane R Mould and Richard Neil Upton. Basic concepts in population modeling, simulation, and model-based drug development—part 2: introduction to pharmacokinetic modeling methods. *CPT: pharmacometrics & systems pharmacology*, 2(4):1–14, 2013.

Arkadi Nemirovski. Lectures on robust convex optimization. Lecture notes, Georgia Institute of Technology, 2012. URL http://www2.isye.gatech.edu/~nemirovs/RO_LN.pdf. Accessed: 2025-05-11.

J Nicholas Neuberger, Alen Alexanderian, and Bart van Bloemen Waanders. Goal oriented optimal design of infinite-dimensional Bayesian inverse problems using quadratic approximations. *arXiv preprint arXiv:2411.07532*, 2024.

David C Norris. How large must a dose-optimization trial be? *CPT: Pharmacometrics & Systems Pharmacology*, 12(11):1777–1783, 2023.

Cameron Nowzari, Victor M. Preciado, and George J. Pappas. Analysis and control of epidemics: A survey of spreading processes on complex networks. *IEEE Control Systems Magazine*, 36(1): 26–46, 2016.

Brendan O'Donoghue, Eric Chu, Neal Parikh, and Stephen Boyd. SCS: Splitting conic solver, version 3.2.7. https://github.com/cvxgrp/scs, November 2023.

Roy Ofir, Francesco Bullo, and Michael Margaliot. Minimum effort decentralized control design for contracting network systems. *IEEE Control Systems Letters*, 6:2731–2736, 2022.

Rafael Orozco, Felix J Herrmann, and Peng Chen. Probabilistic Bayesian optimal experimental design using conditional normalizing flows. *arXiv preprint arXiv:2402.18337*, 2024.

Philip E. Paré, Carolyn L. Beck, and Tamer Başar. Modeling, estimation, and analysis of epidemics over networks: An overview. *Annual Reviews in Control*, 50:345–360, 2020.

Victor M. Preciado, Michael Zargham, Chinwendu Enyioha, Ali Jadbabaie, and George J. Pappas. Optimal resource allocation for network protection against spreading processes. *IEEE Transactions on Control of Network Systems*, 1(1):99–108, 2014.

Tom Rainforth, Rob Cornish, Hongseok Yang, Andrew Warrington, and Frank Wood. On nesting Monte Carlo estimators. In Jennifer Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 4267–4276. PMLR, 10–15 Jul 2018. URL `https://proceedings.mlr.press/v80/rainforth18a.html`.

Tom Rainforth, Adam Foster, Desi R Ivanova, and Freddie Bickford Smith. Modern Bayesian experimental design. *Statistical Science*, 39(1):100–114, 2024.

Elizabeth G Ryan, Christopher C Drovandi, James M McGree, and Anthony N Pettitt. A review of modern computational algorithms for Bayesian optimal design. *International Statistical Review*, 84(1):128–154, 2016.

Rui Silva, Helena Colom, Anabela Almeida, Joana Bicker, Andreia Carona, Ana Silva, Francisco Sales, Isabel Santana, Amílcar Falcão, and Ana Fortuna. A new population pharmacokinetic model for dosing optimization of zonisamide in patients with refractory epilepsy. *European Journal of Pharmaceutical Sciences*, 207:107023, 2025.

Keith D. Smith and Francesco Bullo. Convex optimization of the basic reproduction number. *IEEE Transactions on Automatic Control*, 68(7):4398–4404, 2023.

Dominik Strutz and Andrew Curtis. Variational Bayesian experimental design for geophysical applications: seismic source location, amplitude versus offset inversion, and estimating co2 saturations in a subsurface reservoir. *Geophysical Journal International*, 236(3):1309–1331, 2024.

Mahtab Talaei, Apostolos I Rikos, Alex Olshevsky, Laura F White, and Ioannis Ch Paschalidis. Network-based epidemic control through optimal travel and quarantine management. *arXiv preprint arXiv:2407.19133*, 2024.

Panagiotis Tigas, Yashas Annadani, Andrew Jesson, Bernhard Schölkopf, Yarin Gal, and Stefan Bauer. Interventions, where and how? experimental design for causal models at scale. *Advances in neural information processing systems*, 35:24130–24143, 2022.

Zuzana Vitková, Martin Dodek, Eva Miklovičová, Jarmila Pavlovičová, Andrej Babinec, and Anton Vitko. Robust control of repeated drug administration with variable doses based on uncertain mathematical model. *Bioengineering*, 10(8):921, 2023.

Alžběta Zavřelová, Sara Merdita, Pavel Žák, Jakub Radocha, Benjamin Víšek, Miriam Lánská, Jana Maláková, Pavel Michálek, Ondřej Slanař, and Martin Šíma. Population pharmacokinetic model-based optimization of linezolid dosing in hematooncological patients with suspected or proven gram-positive sepsis. *Clinical and Translational Science*, 18(9):e70346, 2025.

Shijie Zhong, Wanggang Shen, Tommie Catanach, and Xun Huan. Goal-oriented Bayesian optimal experimental design for nonlinear models using markov chain monte carlo. *arXiv preprint arXiv:2403.18072*, 2024.

Huixi Zou, Parikshit Banerjee, Sharon Shui Yee Leung, and Xiaoyu Yan. Application of pharmacokinetic-pharmacodynamic modeling in drug delivery: development and challenges. *Frontiers in pharmacology*, 11:997, 2020.

A Related Work

Goal-oriented Bayesian optimal experimental design Recent advances in BOED have shifted focus from parameter estimation to optimizing for specific quantities of interest (QoIs) — measurable outcomes that directly impact decision-making. For linear models, Attia et al. (2018) established the framework for goal-oriented optimal design of experiments (GOODE) that simplifies computational evaluation for experimental design. Building on this work, Neuberger et al. (2024) introduced a "Gq-optimality" criterion based on quadratic approximation of goal functionals for PDE-governed linear inverse problems. Additionally, for Bayesian linear inverse problems, Madhavan et al. (2025) developed a control-oriented approach that connects optimal control and sensor placement while prioritizing uncertainty reduction in controlled state variables. The linearity in these models makes the problems mathematically tractable and computationally efficient to solve. However, many real-world systems, including epidemic models, exhibit significant nonlinearities that require more sophisticated approaches.

For handling non-linear models, Zhong et al. (2024) created a computational framework using nested Monte Carlo estimators, Markov chain Monte Carlo (MCMC), kernel density estimation, and Bayesian optimization to address both non-linear observation models and prediction models. Similarly, Bickford Smith et al. (2023) proposed the expected predictive information gain (EPIG), an acquisition function that measures information gain in the space of predictions rather than parameters. Taking a different approach, Huang et al. (2024) introduced a decision-aware framework with a transformer neural decision process that simultaneously generates experimental designs and infers decisions in a unified workflow. For causal discovery problems, Tigas et al. (2022) developed methods to optimize intervention timing for large nonlinear structural causal models. Collectively, these works represent a paradigm shift toward experimental designs that optimize directly for decision-relevant outcomes rather than intermediate parameter estimates.

Bayesian decision theory Bayesian decision theory, which applies observed data to update posterior distributions for optimal decision-making, was formalized in Chaloner & Verdinelli (1995). Building on this foundation, Lacoste–Julien et al. (2011) developed a method that calibrates approximate inference techniques according to specific decision tasks using the Expectation-Maximization algorithm. For modern machine learning applications, Krishnan & Tickoo (2020b) introduced a differentiable approach that balances accuracy against uncertainty calibration, enabling models to learn well-calibrated uncertainties while improving performance. Addressing computational efficiency challenges, Gordon et al. (2018) developed a framework that uses few-shot learning to simplify posterior inference of task-specific parameters, eliminating the need for gradient-based optimization during testing. These advances have progressively made Bayesian decision-making more practical for complex problems with computational constraints.

Robust decision-making With the growing interest in goal-oriented BOED, robust decision-making has been studied in many application domains. For example, compartmental network-based approaches (e.g., SIQR model) are widely adopted in epidemic management. Two main control strategies dominate current research: optimal control to minimize infection rates (Lee et al., 2010; Hayhoe et al., 2021; Khanafer & Başar, 2014; Liu & Buss, 2020; Bock & Jayathunga, 2018) and spectral optimization for resource allocation (Hota et al., 2021; Mai et al., 2018; Smith & Bullo, 2023; Preciado et al., 2014; Enyioha et al., 2015). A significant challenge with these approaches is their computational complexity, as many of the underlying problems are NP-complete or NP-hard (Mieghem et al., 2011). In a parallel vein, PK models play a crucial role in optimizing drug dosing and improving patient outcomes by quantitatively linking individual variability to clinical efficacy and safety (Agema et al., 2025; Lai et al., 2022). These models help guide dose selection and treatment personalization, especially under uncertainty in drug absorption, metabolism, and patient response (Zavřelová et al., 2025; Norris, 2023).

B Quarantine Optimization and Lagrangian Formulation

The goal of our quarantine strategy is to minimize implementation costs while effectively controlling disease stability. To achieve this, we define the following objective function:

$$J(q^a, q^s) = \frac{z^a}{1 - q^a} + \frac{z^s}{1 - q^s},$$

and we solve the following minimization problem:

$$\min_{q^a, q^s} J(q^a, q^s)$$

subject to the constraints:

$$\begin{aligned} \lambda_{\max}(M(t_0, \boldsymbol{\beta}_i, \boldsymbol{\gamma}_i)) \leq -\alpha \quad \text{for} \quad i = 1, \ldots, N, \\ 0 \leq q^a \leq 1, \\ 0 \leq q^s \leq 1, \end{aligned}$$

where $\boldsymbol{\beta}_i$ and $\boldsymbol{\gamma}_i$ are sampled from the posterior distribution $p(\boldsymbol{\beta}, \boldsymbol{\gamma}|y, \xi)$, and N denotes the number of samples.

To address this constrained optimization problem, we introduce the Lagrangian:

$$\begin{aligned} \mathcal{L} = \frac{z^a}{1-q^a} + \frac{z^s}{1-q^s} + \sum_{i=1}^{N} \lambda_i (\lambda_{\max}(M(t_0, \boldsymbol{\beta}_i, \boldsymbol{\gamma}_i)) + \alpha) \\ - \mu_1 q^a + \mu_2 (q^a - 1) - \mu_3 q^s + \mu_4 (q^s - 1), \end{aligned}$$

where λ_i are Lagrange multipliers associated with the eigenvalue constraints, and μ_i are multipliers for the box constraints on q^a and q^s. The optimal solution must satisfy the Karush-Kuhn-Tucker (KKT) conditions, which we outline below.

B.1 KKT CONDITIONS

Stationarity The stationarity conditions are derived by taking partial derivatives of the Lagrangian:

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial q^a} = \frac{z^a}{(1-q^a)^2} + \sum_{i=1}^{N} \lambda_i \frac{\partial \lambda_{\max}(M(t_0, \boldsymbol{\beta}_i, \boldsymbol{\gamma}_i))}{\partial q^a} + \mu_2 - \mu_1 = 0, \\ \frac{\partial \mathcal{L}}{\partial q^s} = \frac{z^s}{(1-q^s)^2} + \sum_{i=1}^{N} \lambda_i \frac{\partial \lambda_{\max}(M(t_0, \boldsymbol{\beta}_i, \boldsymbol{\gamma}_i))}{\partial q^s} + \mu_4 - \mu_3 = 0. \end{aligned}$$

Primal feasibility The primal feasibility conditions ensure the constraints hold:

$$\begin{aligned} \lambda_{\max}(M(t_0, \boldsymbol{\beta}_i, \boldsymbol{\gamma}_i)) \leq -\alpha \quad \text{for} \quad i = 1, \ldots, N, \\ 0 \leq q^a \leq 1, \\ 0 \leq q^s \leq 1. \end{aligned}$$

Dual feasibility The Lagrange multipliers must be non-negative:

$$\lambda_i \geq 0 \quad \text{for} \quad i = 1, \ldots, N, \quad \text{and} \quad \mu_1, \mu_2, \mu_3, \mu_4 \geq 0.$$

Complementary slackness The complementary slackness conditions are:

$$\begin{aligned} \lambda_i \left(\lambda_{\max}(M(t_0, \boldsymbol{\beta}_i, \boldsymbol{\gamma}_i)) + \alpha\right) = 0 \quad \text{for} \quad i = 1, \ldots, N, \\ \mu_1 q^a = 0, \\ \mu_2 (1 - q^a) = 0, \\ \mu_3 q^s = 0, \\ \mu_4 (1 - q^s) = 0. \end{aligned}$$

Assuming an interior solution, the multipliers for the box constraints at the optimal point become $\mu_i = 0$. The KKT conditions simplify to:

$$\frac{z^a}{(1-\hat{q}^a)^2} + \sum_{i=1}^{N} \hat{\lambda}_i \frac{\partial \lambda_{\max}(M(t_0, \boldsymbol{\beta}_i, \boldsymbol{\gamma}_i))}{\partial q^a} = 0,$$

$$\frac{z^s}{(1-\hat{q}^s)^2} + \sum_{i=1}^{N} \hat{\lambda}_i \frac{\partial \lambda_{\max}(M(t_0, \boldsymbol{\beta}_i, \boldsymbol{\gamma}_i))}{\partial q^s} = 0,$$

$$\lambda_{\max}(M(t_0, \boldsymbol{\beta}_{i^*}, \boldsymbol{\gamma}_{i^*})) = -\alpha \quad \text{for specific} \quad i^*,$$

where $\hat{\lambda}_i$ represents for the Lagrangians at the optimal point, and i^* indicates the specific boundary condition corresponding to the optimal point.

We compute the optimal values $\hat{q}^a$, $\hat{q}^s$, and $\hat{\lambda}_{i^*}$ using convex optimization via semidefinite programming. This is feasible because J is a convex function with respect to q^a and q^s, and the largest eigenvalue constraint can be reformulated as a set of linear matrix inequalities. For implementation details, see Appendix C.

B.2 Derivative of the optimum cost J^* with respect to model parameters

We can compute the gradient of the optimum cost J^* with respect to the model parameters $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$. Using the envelope theorem, the derivative with respect to β_i^a is:

$$\frac{\partial J^*}{\partial \beta_i^a} = \sum_{j=1}^{N} \frac{\partial \hat{\lambda}_j (\lambda_{\max}(M(t_0), \boldsymbol{\beta}_j, \boldsymbol{\gamma}_j) + \alpha)}{\partial \beta_i^a} = \hat{\lambda}_i \cdot \frac{\partial \lambda_{\max}(M(t_0, \boldsymbol{\beta}_i, \boldsymbol{\gamma}_i))}{\partial \beta_i^a}.$$

Similarly, the derivatives with respect to other parameters are:

$$\frac{\partial J^*}{\partial \gamma_i^s} = \hat{\lambda}_i \cdot \frac{\partial \lambda_{\max}(M(t_0, \boldsymbol{\beta}_i, \boldsymbol{\gamma}_i))}{\partial \gamma_i^s},$$

$$\frac{\partial J^*}{\partial \gamma_i^a} = \hat{\lambda}_i \cdot \frac{\partial \lambda_{\max}(M(t_0, \boldsymbol{\beta}_i, \boldsymbol{\gamma}_i))}{\partial \gamma_i^a},$$

$$\frac{\partial J^*}{\partial \beta_i^s} = \hat{\lambda}_i \cdot \frac{\partial \lambda_{\max}(M(t_0, \boldsymbol{\beta}_i, \boldsymbol{\gamma}_i))}{\partial \beta_i^s}.$$

Since the optimal solution often lies on specific boundaries where $\hat{\lambda}_{i^*} \neq 0$ and $\hat{\lambda}_i = 0$ for $i \neq i^*$, the gradient depends only on the samples directly influencing the solution.

B.3 Derivative of the constraint $h(y, \xi)$ with respect to experimental design ξ

The constraint term is defined as:

$$h(y, \xi) = \sum_{i=1}^{N} \lambda_i \left(\lambda_{\max}(M(t_0, \boldsymbol{\beta}_i, \boldsymbol{\gamma}_i)) + \alpha\right) - \mu_1 q^a + \mu_2 (q^a - 1) - \mu_3 q^s + \mu_4 (q^s - 1).$$

Its gradient with respect to ξ at the optimal point is:

$$\nabla_\xi h(y, \xi) = \sum_{i=1}^{N} \hat{\lambda}_i \left(\frac{\partial \lambda_{\max}}{\partial \beta_i^a} \cdot \nabla_\xi \beta_i^a + \frac{\partial \lambda_{\max}}{\partial \beta_i^s} \cdot \nabla_\xi \beta_i^s + \frac{\partial \lambda_{\max}}{\partial \gamma_i^a} \cdot \nabla_\xi \gamma_i^a + \frac{\partial \lambda_{\max}}{\partial \gamma_i^s} \cdot \nabla_\xi \gamma_i^s \right). \quad (13)$$

The partial derivative $\frac{\partial \lambda_{\max}}{\partial \beta_i^a}$ is given by Talaei et al. (2024) as:

$$\frac{\partial \lambda_{\max}}{\partial \beta_i^a} = \frac{v_{\max}^T \left(\frac{\partial M(t_0, \boldsymbol{\beta}_i)}{\partial \beta_i^a} \right) u_{\max}}{v_{\max}^T u_{\max}},$$

where $v_{\max}$ and $u_{\max}$ are the left and right eigenvectors of the largest eigenvalue, respectively. The terms $\nabla_\xi \beta_i^a$, $\nabla_\xi \beta_i^s$, $\nabla_\xi \gamma_i^a$, and $\nabla_\xi \gamma_i^s$ are computed via automatic differentiation from the variational network.

B.4 MARGINAL LIKELIHOOD GRADIENT

To compute the log-likelihood gradient with respect to ξ in **??**, we assume a Poisson observation model (suitable for count data) with rate parameter $\lambda = 0.95 \cdot y_{\text{true}}(\xi)$. The gradient is:

$$\frac{\partial}{\partial \xi} \log p(y_{\text{obs}}|\xi) = \left(\frac{y_{\text{obs}}}{0.95 \cdot y_{\text{true}}(\xi)} - 1 \right) \cdot 0.95 \cdot \frac{\partial y_{\text{true}}(\xi)}{\partial \xi}.$$

This expression helps quantify how changes in ξ affect the likelihood of the observed data. The term $\frac{\partial y_{\text{true}}(\xi)}{\partial \xi}$ can be approximated using finite difference methods, such as the central difference method.

C OPTIMIZATION FOR OPTIMAL CONTROL AND NEURAL NETWORK FOR VARIATIONAL INFERENCE DETAILS

We utilize SCS O'Donoghue et al. (2023) and MOSEK ApS (2025) to express the semi-definite programming problem and use MOSEK's implementation of the interior point method with default settings for optimization

Architecture. Let $\xi \in \Xi$ denote the design and $y \in \mathcal{Y}$ the observation. Both are mapped into a shared latent space using linear tokenizers that produce M tokens of width d:

$$E_\xi : \mathbb{R}^{D_\xi} \to \mathbb{R}^{M \times d}, \quad E_y : \mathbb{R}^{D_y} \to \mathbb{R}^{M \times d}, \qquad Z_\xi = E_\xi(\xi),\ Z_y = E_y(y),$$

where in PK we set $D_\xi = 1$, $D_y = 1$, $M = 8$, and $d = 64$. Queries and keys are linear projections of the token sequences,

$$q_j = W_q\, z_{y,j} \in \mathbb{R}^d, \qquad k_i = W_k\, z_{\xi,i} \in \mathbb{R}^d,$$

and values fuse per-token key and query latents via a small MLP,

$$v_i = \Psi\big([\, z_{\xi,i},\, z_{y,i}\,]\big) \in \mathbb{R}^d, \qquad i, j \in \{1, \dots, M\}.$$

Single-head dot-product cross-attention over tokens is

$$a_{ji} = \frac{\exp\big(q_j^\top k_i / \sqrt{d}\big)}{\sum_{i'=1}^{M} \exp\big(q_j^\top k_{i'} / \sqrt{d}\big)}, \qquad c_j = \sum_{i=1}^{M} a_{ji}\, v_i \in \mathbb{R}^d.$$

We mean-pool the query contexts and pass through a light trunk MLP:

$$s = \frac{1}{M} \sum_{j=1}^{M} c_j, \qquad h = \mathrm{MLP}_{\mathrm{trunk}}(s) \in \mathbb{R}^d.$$

Two heads with skip connections produce log-space parameters for a diagonal LogNormal posterior,

$$\tilde{\mu} = W_{\ell,2}\, \phi(W_{\ell,1} h) + W_{\ell,\mathrm{skip}} h,$$
$$\tilde{\sigma} = W_{s,2}\, \phi(W_{s,1} h) + W_{s,\mathrm{skip}} h,$$

which we bound elementwise:

$$\mu \;=\; \mu_0 \;+\; \Delta_{\max}\, \tanh(\tilde{\mu}), \qquad \sigma \;=\; \sigma_{\min} \;+\; (\sigma_{\max} - \sigma_{\min})\, \sigma\big(\tilde{\sigma}\big).$$

Here ϕ is GELU, μ_0 is the prior log-mean (used as a residual center), $\Delta_{\max}$ bounds deviations, and $0 < \sigma_{\min} < \sigma_{\max}$ bound the log-space standard deviations.

Training Configuration: GPU: NVIDIA A100.

D MODELS

D.1 OPTIMAL CONTROL FOR EPIDEMIOLOGY

Building on the BOED framework, we consider robust epidemiology control as an example. Our method leverages the framework established by Talaei et al. (2024). The epidemiology model is governed by a SIQR spread disease network, where state variables represent different compartments of the population: susceptible (s), asymptomatic infected (x^a), symptomatic infected (x^s) and recovered (h).

$$\begin{pmatrix} \dot{s} \\ \dot{x}^a \\ \dot{x}^s \\ \dot{h} \end{pmatrix} = \begin{pmatrix} 0 & -\beta^a s & -\beta^s s & 0 \\ 0 & \beta^a s - \epsilon - \gamma^a - q^a & \beta^s s & 0 \\ 0 & \epsilon & -\gamma^s - q^s & 0 \\ 0 & \gamma^a & \gamma^s & 0 \end{pmatrix} \begin{pmatrix} s \\ x^a \\ x^s \\ h \end{pmatrix}. \tag{14}$$

Here, β^a and β^s are transmission rates for asymptomatic and symptomatic cases, ϵ is the rate at which asymptomatic cases develop symptoms, γ^a and γ^s are recovery rates for asymptomatic and symptomatic cases, and q^a and q^s are quarantine rates for asymptomatic and symptomatic individuals.

Following Ma et al. (2023), we decouple the dynamics of $\dot{x}$ from $\dot{s}$ and $\dot{h}$, allowing us to focus on the matrix $M(t_0)$ which captures the essential infection dynamics at initial time t_0:

$$M(t_0) = \begin{pmatrix} \beta^a s(t_0) - \epsilon - \gamma^a - q^a & \beta^s s(t_0) \\ \epsilon & -\gamma^s - q^s \end{pmatrix} \tag{15}$$

To optimize the quarantine strategy as described in Talaei et al. (2024), we utilize an objective function that minimizes economic costs:

$$\min_{q^a, q^s} J(q^a, q^s) = \frac{z^a}{1 - q^a} + \frac{z^s}{1 - q^s} \tag{16}$$

$$\text{s.t. } \lambda_{\max}(M(t_0)) \leq -\alpha, \tag{17}$$

$$0 \leq q^a \leq 1, \tag{18}$$

$$0 \leq q^s \leq 1, \tag{19}$$

where z^a represents the economic cost for asymptomatic quarantine, z^s represents the economic cost for symptomatic quarantine, and $\alpha > 0$ is a constraint ensuring the stability of the system.

The detailed solution for this minimization problem is provided in Appendix B. This approach allows us to determine optimal quarantine rates without explicitly integrating the SIQR differential equations. Instead, by analyzing the eigenvalues of the system and applying convex optimization, we can efficiently identify the optimal quarantine strategy that minimizes economic costs.

D.2 OPTIMAL CONTROL FOR PHARMACOKINETIC MODEL

We further consider PK model as another example. The concentration at time t for a drug administered orally can be modeled using the Bateman function (Bateman, 1910):

$$y(t) = \frac{D \cdot k_a}{V \cdot (k_a - k_e)} \left(e^{-k_e \cdot t} - e^{-k_a \cdot t}\right) (1 + \epsilon_{\text{mult}}) + \epsilon_{\text{add}},$$

where V is the volume of distribution, k_a is the absorption rate constant, k_e is the elimination rate constant, D is the dose administered, ϵ_{mult} is a multiplicative error term, and ϵ_{add} is additive error term. This formulation captures the dynamics of drug absorption and elimination.

Given the drug's potential toxicity, dosing should maintain systemic exposure within the therapeutic window—avoiding toxic concentrations while not falling below the minimum effective concentration. We can calculate the time at which the maximum drug concentration occurs, denoted $t_{\max}$, can be found by setting the derivative $\partial y / \partial t = 0$, which yields:

$$t_{\max} = \frac{\ln(k_a / k_e)}{k_a - k_e}.$$

The maximum concentration, $C_{\max}$, is obtained by evaluating $y(t)$ at $t_{\max}$:

$$y_{\max} = \frac{D}{V}\left(\frac{k_e}{k_a}\right)^{\frac{k_e}{-k_e+k_a}}(1+\epsilon_{\text{mult}}) + \epsilon_{\text{add}}.$$

For cumulative exposure, the area under the concentration curve (AUC) is given by:

$$\text{AUC} = \int_0^\infty y(t)dt = \frac{D}{V \cdot k_e},$$

assuming complete absorption.

We define a convex cost function $J(q)$ (e.g., $J(q) = c\,q$ to discourage high dosing). The constrained problem can be written in the following forms,

$$\begin{aligned} \min_{0\le q\le 1} \quad & J(q) \\ \text{s.t.} \quad & y_{\max}(q,\boldsymbol{\theta}) \le y_{\text{thresh}}, \\ & \text{AUC}(q,\boldsymbol{\theta}) \;\ge\; \text{AUC}_{\min}. \end{aligned} \tag{20}$$

E Robust Decision Making

E.1 SIQR model

We develop an integrated framework, GoBOED, for epidemic management by bringing together BOED (introduced in eq. (2)) and optimal control through convex optimization (described in eq. (16)). Our primary objective is to identify an optimal experimental design ξ^* that minimizes the expected controlled economic cost, where the expectation is taken over posterior distributions of the parameters $\boldsymbol{\beta} = (\beta^a, \beta^s)$ and $\boldsymbol{\gamma} = (\gamma^a, \gamma^s)$. This approach allows us to update our beliefs with new observations and improve decision-making, following the methodology of Chaloner & Verdinelli (1995).

For a given experimental design ξ and observed data y, we formulate the optimization problem as:

$$\min_{q^a,q^s} \mathbb{E}_{p(\boldsymbol{\beta},\boldsymbol{\gamma}|y,\xi)}\left[J(q^a,q^s)\right] \tag{21}$$

$$\begin{aligned} \text{s.t. } \lambda_{\max}\left(\mathbb{E}_{p(\boldsymbol{\beta},\boldsymbol{\gamma}|y,\xi)}\left[M(t_0,\boldsymbol{\beta},\boldsymbol{\gamma})\right]\right) &\le -\alpha, \\ 0 \le q^a &\le 1, \\ 0 \le q^s &\le 1. \end{aligned} \tag{22}$$

Here, $J(q^a, q^s)$ represents the economic cost, which does not directly depend on $\boldsymbol{\beta}$ or $\boldsymbol{\gamma}$ but is constrained by the constraints in eq. (22). Thus, the optimization problem simplifies to minimizing $J(q^a, q^s)$ subject to the constraints, where the key challenge lies in evaluating the eigenvalue constraint. This constraint requires computing the expectation of the matrix $M(t_0, \boldsymbol{\beta}, \boldsymbol{\gamma})$ over the posterior distribution $p(\boldsymbol{\beta}, \boldsymbol{\gamma}|y, \xi)$.

To approximate the posterior distribution, we employ stochastic variational inference. Specifically, we approximate the true posterior $p(\boldsymbol{\beta}, \boldsymbol{\gamma}|y, \xi)$ using a variational distribution $q_\phi(\boldsymbol{\beta}, \boldsymbol{\gamma}|y, \xi)$, where ϕ denotes the variational parameters. The variational distribution is optimized by maximizing the evidence lower bound (ELBO), which can be expressed as:

$$\mathcal{L}_{VI}(\phi; y) = \mathbb{E}_{q_\phi(\boldsymbol{\beta},\boldsymbol{\gamma}|y,\xi)}\left[\log p(y|\boldsymbol{\beta},\boldsymbol{\gamma},\xi)\right] - D_{KL}\left(q_\phi(\boldsymbol{\beta},\boldsymbol{\gamma}|y,\xi)\|p(\boldsymbol{\beta},\boldsymbol{\gamma})\right). \tag{23}$$

Once we have obtained the variational posterior $q_\phi(\boldsymbol{\beta}, \boldsymbol{\gamma}|y, \xi)$, we use it to estimate the expectation in the eigenvalue constraint. Since q_ϕ may not exactly match the true posterior, we apply importance sampling to correct for the discrepancy. For any function $f(\boldsymbol{\beta}, \boldsymbol{\gamma})$, the expectation under the true posterior can be estimated as

$$\mathbb{E}_{p(\boldsymbol{\beta},\boldsymbol{\gamma}|y,\xi)}[f(\boldsymbol{\beta},\boldsymbol{\gamma})] \approx \frac{\sum_{i=1}^N f(\boldsymbol{\beta}_i,\boldsymbol{\gamma}_i)w(\boldsymbol{\beta}_i,\boldsymbol{\gamma}_i)}{\sum_{i=1}^N w(\boldsymbol{\beta}_i,\boldsymbol{\gamma}_i)}, \tag{24}$$

where $(\boldsymbol{\beta}_i, \boldsymbol{\gamma}_i)$ are samples from $q_\phi(\boldsymbol{\beta}, \boldsymbol{\gamma}|y, \xi)$, and $w(\boldsymbol{\beta}_i, \boldsymbol{\gamma}_i) = \frac{p(y|\boldsymbol{\beta}_i, \boldsymbol{\gamma}_i, \xi)p(\boldsymbol{\beta}_i, \boldsymbol{\gamma}_i)}{q_\phi(\boldsymbol{\beta}_i, \boldsymbol{\gamma}_i|y, \xi)}$ are the importance weights.

In particular, for the eigenvalue constraint, we compute $\mathbb{E}_{p(\boldsymbol{\beta}, \boldsymbol{\gamma}|y, \xi)}[M(t_0, \boldsymbol{\beta}, \boldsymbol{\gamma})]$ using the above formula with $f(\boldsymbol{\beta}, \boldsymbol{\gamma}) = M(t_0, \boldsymbol{\beta}, \boldsymbol{\gamma})$. This allows us to evaluate the constraint and solve the optimization problem for each ξ and y. Thus, our framework integrates BOED with optimal control under posterior distribution uncertainty, leveraging variational inference and importance sampling.

E.1.1 Alternative formulation using chance constraints

To address the uncertainty in parameters $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$ robustly, we propose an alternative formulation using chance constraints. Chance constraints are a powerful tool in optimization under uncertainty, ensuring that critical conditions hold with a specified probability (Charnes & Cooper, 1959; Miller & Wagner, 1965). This approach simplifies our problem by focusing on the eigenvalue constraint while managing parameter uncertainty effectively.

Given that the objective function is independent of $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$, we can reformulate the optimization problem to concentrate on the maximum eigenvalue conditions with the posterior samples. Directly imposing eigenvalue constraints for each sample would be overly restrictive. Instead, we introduce a chance constraint to ensure that the maximum eigenvalue of the matrix $M(t_0, \boldsymbol{\beta}, \boldsymbol{\gamma})$ satisfies a stability condition with high probability.

The alternative formulation, denoted $\mathcal{L}_{\text{OC-CC}}$, is defined as:

$$\min_{q^a, q^s} J(q^a, q^s) \tag{25}$$

$$\text{s.t. } P\left(\lambda_{\max}(M(t_0, \boldsymbol{\beta}, \boldsymbol{\gamma})) \leq -\alpha \middle| \, y, \xi\right) \geq \eta, \tag{26}$$

$$0 \leq q^a \leq 1, \tag{27}$$

$$0 \leq q^s \leq 1, \tag{28}$$

where $(\boldsymbol{\beta}, \boldsymbol{\gamma})$ comes from the posterior distribution $p(\boldsymbol{\beta}, \boldsymbol{\gamma}|y, \xi)$, and $\eta \in (0, 1)$ is the desired probability level.

To evaluate the chance constraint eq. (26) efficiently, we employ importance sampling with variational inference. We set $f(\boldsymbol{\beta}, \boldsymbol{\gamma})$ in eq. (24) to be the indicator function $I(\lambda_{\max}(M(t_0, \boldsymbol{\beta}, \boldsymbol{\gamma})) \leq -\alpha)$, allowing us to evaluate $P\left(\lambda_{\max}(M(t_0, \boldsymbol{\beta}, \boldsymbol{\gamma})) \leq -\alpha \middle| \, y, \xi\right)$. This formulation enables us to directly assess the constraint and optimize the economic cost accordingly. Without this approach, samples far from the mean would make the constraint evaluation numerically unstable and potentially cause the optimization to diverge. The procedure is implemented as follows:

1. Draw N samples $\{(\boldsymbol{\beta}_i, \boldsymbol{\gamma}_i)\}_{i=1}^N$ from $q_\phi(\boldsymbol{\beta}, \boldsymbol{\gamma}|y, \xi)$.
2. Compute importance weights: $w_i = \frac{p(\boldsymbol{\beta}_i, \boldsymbol{\gamma}_i)p(y|\boldsymbol{\beta}_i, \boldsymbol{\gamma}_i, \xi)}{q(\boldsymbol{\beta}_i, \boldsymbol{\gamma}_i|y, \xi)}$.
3. Normalize the weights: $\tilde{w}_i = \frac{w_i}{\sum_{j=1}^N w_j}$.
4. Estimate the conditional probability for a given y and ξ:

$$P(\lambda_{\max}(M(t_0, \boldsymbol{\beta}, \boldsymbol{\gamma})) \leq -\alpha|y, \xi) \approx \sum_{i=1}^N \tilde{w}_i I(\lambda_{\max}(M(t_0, \boldsymbol{\beta}_i, \boldsymbol{\gamma}_i)) \leq -\alpha),$$

   where $I(\cdot)$ is the indicator function.

With the selected posterior samples, we can formulate constraints on the largest eigenvalue and solve the resulting optimization problem using convex optimization in a semidefinite programming framework.

E.1.2 CVaR-based robust control

As an alternative to the scenario-based chance constraint in Appendix E.1.1, we control the *tail* of constraint violations using Conditional Value-at-Risk (CVaR).

Let the feasibility event be

$$P\big(\lambda_{\max}\big(M(t_0, \boldsymbol{\beta}, \boldsymbol{\gamma}, q^a, q^s)\big) \le -\alpha \,\big|\, y, \xi\big) \;\ge\; \eta.$$

For samples $\{(\boldsymbol{\beta}_i, \boldsymbol{\gamma}_i)\}_{i=1}^N$ (posterior or variational), define the per-sample violation

$$g_i(q^a, q^s) \;:=\; \lambda_{\max}\big(M(t_0, \boldsymbol{\beta}_i, \boldsymbol{\gamma}_i, q^a, q^s)\big) \;+\; \alpha.$$

We enforce $\mathrm{CVaR}_\eta\big(g(\Xi, q)\big) \le 0$ using the Rockafellar–Uryasev sample-average form with optional normalized weights $\bar{w}_i$ (default $\bar{w}_i = \frac{1}{N}$):

$$s_i \;\ge\; g_i(q^a, q^s) - \tau, \quad i = 1, \dots, N, \qquad \tau + \frac{1}{1-\eta} \sum_{i=1}^N \bar{w}_i \, s_i \;\le\; 0, \tag{29}$$

with decision variables $\tau \in \mathbb{R}$ and $s_i \ge 0$. At optimality, $s_i = (g_i(q^a, q^s) - \tau)_+$, so only the upper tail beyond the η-quantile contributes.

Our robust controls are then obtained by solving the convex program

$$\begin{aligned} \min_{q^a, q^s,\, \tau,\, \{s_i\}} \quad & J(q^a, q^s) \\ \text{s.t.} \quad & s_i \;\ge\; g_i(q^a, q^s) - \tau, \;\; s_i \ge 0, \;\; i = 1, \dots, N, \\ & \tau + \frac{1}{1-\eta} \sum_{i=1}^N \bar{w}_i \, s_i \;\le\; 0, \\ & 0 \le q^a \le 1, \;\; 0 \le q^s \le 1, \end{aligned} \tag{30}$$

where $J(q^a, q^s)$ is the same convex objective as in Appendix E.1.1. The specific form of g_i is model dependent; in our case it is implemented through a convex epigraph for the spectral violation, so eq. (30) remains a tractable convex program.

E.2 PK MODEL

Parameter influence. $C_{\max} \propto D/V$ and increases with k_a (faster absorption, smaller $t_{\max}$), decreases with k_e (faster elimination), and is sensitive to the ratio k_a/k_e (flip–flop when $k_a < k_e$).

We follow the design–control split used in the SIQR GoBOED paper. Let ξ: e.g., blood sampling schedule and/or formulation choice, used to learn $\boldsymbol{\theta} = (k_a, k_e, V)$ via posterior $p(\boldsymbol{\theta} \mid y, \xi)$. $q \in [0, 1]$: a dose fraction, with $D(q) = q\, D_0$ (replace by any convex mapping as needed). All risk and exposure quantities below depend on q through $D(q)$ and on $\boldsymbol{\theta}$.

E.2.1 POSTERIOR EXPECTATIONS VIA VI + IMPORTANCE WEIGHTING

Given a variational posterior $q_\phi(\boldsymbol{\theta} \mid y, \xi)$, expectations under the true posterior are estimated by importance weighting:

$$\mathbb{E}_{p(\boldsymbol{\theta}|y,\xi)}[f(\boldsymbol{\theta})] \;\approx\; \sum_{i=1}^N \tilde{w}_i \, f(\boldsymbol{\theta}_i), \qquad \tilde{w}_i = \frac{w_i}{\sum_{j=1}^N w_j}, \quad w_i = \frac{p(y \mid \boldsymbol{\theta}_i, \xi)\, p(\boldsymbol{\theta}_i)}{q_\phi(\boldsymbol{\theta}_i \mid y, \xi)}, \;\; \boldsymbol{\theta}_i \sim q_\phi. \tag{31}$$

E.2.2 CHANCE-CONSTRAINED CONTROL (PK ANALOGUE OF SIQR)

Choose a convex cost $J(q)$ (e.g., $J(q) = c\, q$ to discourage high dosing). The chance-constrained problem mirrors the SIQR formulation:

$$\begin{aligned} \min_{0 \le q \le 1} \quad & J(q) \\ \text{s.t.} \quad & \mathbb{P}(C_{\max}(q, \boldsymbol{\theta}) \le C_{\text{thresh}} \,|\, y, \xi) \;\ge\; \eta, \\ & \mathbb{E}[\mathrm{AUC}(q, \boldsymbol{\theta})] \;\ge\; \mathrm{AUC}_{\min}. \end{aligned} \tag{32}$$

The probability and expectation are evaluated using eq. (31) with $f(\boldsymbol{\theta}) = \mathbf{1}\{C_{\max}(q, \boldsymbol{\theta}) \le C_{\text{thresh}}\}$ and $f(\boldsymbol{\theta}) = \mathrm{AUC}(q, \boldsymbol{\theta})$, respectively.

E.2.3 CVaR-ROBUST CONTROL (EPIGRAPH FORM)

Define per-sample violation (optionally with a nonnegative safety margin α)

$$g_i(q) \;=\; C_{\max}(q, \boldsymbol{\theta}_i) \;-\; C_{\text{thresh}} \;+\; \alpha, \qquad i = 1, \ldots, N, \tag{33}$$

with normalized importance weights $\tilde{w}_i$ from eq. (31). Enforce $\mathrm{CVaR}_\eta(g) \le 0$ via

$$\begin{aligned} \min_{0 \le q \le 1,\ \tau,\ s_i \ge 0} \quad & J(q) \\ \text{s.t.} \quad & s_i \;\ge\; g_i(q) - \tau, \qquad i = 1, \ldots, N, \\ & \tau \;+\; \frac{1}{1-\eta} \sum_{i=1}^{N} \tilde{w}_i \, s_i \;\le\; 0, \\ & \mathbb{E}[\mathrm{AUC}(q, \boldsymbol{\theta})] \;\ge\; \mathrm{AUC}_{\min}. \end{aligned} \tag{34}$$

F NUMERICAL RESULTS

For our implementation of the SIQR model, we adopted the parameterization framework established by Talaei et al. (2024). We specified log-normal prior distributions for both transmission and recovery rates in units of counts per day. Specifically, we assigned log-normal distributions to the transmission rates for asymptomatic individuals (β^a) and symptomatic individuals (β^s), as well as to the recovery rates for asymptomatic individuals (γ^a) and symptomatic individuals (γ^s). The natural logarithm of these distributions have means of $(0.5, 0.8, 0.2, 0.2)$ respectively, with a standard deviation of $(0.5, 0.5, 0.3, 0.3)$ for each parameter. The stability parameter (α) was fixed at 0.05. The economic cost parameters z^a, z^s were set to (0.4, 0.6).

Using these parameters, we solved the SIQR model. The design variable $\xi \in [1, 100]$ represents the observation day. At this observation point, we measure y_{obs}, which consists of counts of asymptomatic and symptomatic infected individuals. We model these observed values using a Poisson distribution where the rate parameter λ equals $0.95 \cdot y_{\text{true}}(\xi)$, with $y_{\text{true}}(\xi)$ being the actual model-predicted values at time ξ. We then train the posterior distribution using the observed data y_{obs} and design variable ξ. When computing the optimization problem for the optimal economic cost, we use ApS (2025) for convex optimization. More details of optimization and training are provided in Appendix C.

For our implementation of the PK model, we adopted the parameterization framework established by Kleinegesse & Gutmann (2021). We specified log-normal prior distributions for k_a, k_e, V. Specifically, we assigned log-normal distributions for the parameters. The natural logarithm of these distributions have means of $(0, \log(0.1), \log(20.0))$ respectively, with a standard deviation of $(0.05, 0.05, 0.05)$ for each parameter.