# Bayesian Risk-Aware CBFs for Discrete-Time Stochastic Systems with Learned Dynamics

Bardh Hoxha, Mitchell Black, Keyvan Majd, Hideki Okamoto, Georgios Fainekos, Danil Prokhorov

*Toyota Motor North America R&D*

***Abstract*—We study safety for stochastic systems under sampled–data control and learned dynamics. We develop Bayesian Risk-Aware Control Barrier Functions (RA-CBFs) for discrete time. First, we give two guarantees for barrier crossing over a finite horizon: (i) a time-uniform martingale bound using Ville's inequality and a predictable variance budget, and (ii) a tighter pathwise bound that recovers the continuous-time RA-CBF margin via DDS time change and the reflection principle. Second, we propagate posterior uncertainty in drift and diffusion into an upper confidence bound (UCB) on the generator and derive a closed-form inter-sample robustness margin. Third, we synthesize a minimally invasive controller via a convex QP that enforces risk and uncertainty thresholds jointly. The result is a high-confidence, finite-horizon safety bound and a practical sampled-data controller that, under the stronger pathwise noise assumption, recovers the continuous-time RA-CBF margin while avoiding supermartingale-based S-CBF conditions.**

***Index Terms*—Risk-Aware Control, Control Barrier Functions, Stochastic Systems, Bayesian Learning, Probabilistic Safety**

## I. Introduction

Ensuring the safety of autonomous systems is paramount as they operate in uncertain environments. Control Barrier Functions (CBFs) have emerged as a powerful tool for enforcing safety constraints, guaranteeing the forward invariance of a constraint set for deterministic systems [1]–[4]. Complicating this guarantee is the fact that real-world models are imperfect, environments are uncertain, and measurements are noisy. Extending these guarantees to real-world scenarios requires addressing stochastic dynamics, model uncertainty, and the discrete-time (DT) nature of digital control.

In recent years, CBFs for safety under uncertainty have seen improvements thanks to synthesis with Gaussian process learning for modeling error [5]–[7] and distributionally robust optimization accounting for noisy measurements [8], [9], unpredictable obstacles [10], or environmental perturbations [11]. Since these approaches reformulate a deterministic barrier function constraint into a chance constraint to be enforced pointwise-in-time, there is typically no guarantee of safety over any measurable time interval. Stochastic CBFs (S-CBFs) were developed in the stochastic differential equation (SDE) modeling paradigm to certify safety probabilistically over a finite time interval [12], [13]. Although S-CBFs rely on a supermartingale condition that can be overly conservative, Risk-Aware CBFs (RA-CBFs) [14] were recently introduced as a less conservative alternative by using tight boundary-crossing probabilities to convert a user-defined risk threshold directly into a required safety margin on the system's generator.

Despite this advance, three critical gaps remain. First, RA-CBFs were formulated in continuous time (CT), whereas physical controllers are digital and operate in discrete time (DT). Second, the framework assumes perfect knowledge of the system dynamics. In practice, drift and diffusion must be learned from data, and this model uncertainty must be formally managed. Third, despite work studying inter-sample safety with CBFs [15]–[17], this remains an open problem for systems modeled as SDEs.

We introduce a Bayesian RA-CBF framework for discrete-time stochastic systems with learned dynamics that provides (i) time-uniform and pathwise-aware DT safety guarantees, (ii) a receding-horizon implementation with UCB-based uncertainty propagation, and (iii) a tractable QP controller with explicit inter-sample robustness. Relative to [14], which treats continuous-time systems with known dynamics, our approach addresses sampled-data control and learned dynamics. Relative to [12], we avoid supermartingale-based conditions by using martingale concentration and a pathwise Brownian-crossing argument. Relative to [18]–[20], we incorporate Bayesian model uncertainty directly into the safety margin rather than assuming known noise statistics.

Closely related is the uncertainty-separated CBF framework of [18], which proposes a heuristic decomposition of epistemic and aleatoric uncertainty, the discrete-time, stochastic CBFs of [19], which uses Freedman's inequality to obtain K-step exit bounds, and the iterative, SMT-verified sampled-data stochastic control synthesis of [20]. In contrast to [19], whose bounds can be tighter than Ville-based results and do not require an upper-bounded barrier function, but whose focus is purely discrete-time, we instead recover the continuous-time RA-CBF margin under a pathwise assumption, handle learned dynamics via a Bayesian generator UCB, and provide inter-sample guarantees. Our work complements [18]–[20] by establishing formal, finite-horizon probabilistic guarantees for computationally efficient, sampled-data control of continuous-time stochastic systems, including explicit treatment of inter-sample behavior and principled posterior-based risk calibration.

Our primary contributions are:

- We establish a two-tier framework for discrete-time safety, providing a robust, time-uniform guarantee via martingale concentration and a tighter pathwise bound that recovers continuous-time performance.
- We introduce a Bayesian RA-CBF that propagates posterior uncertainty over the drift and diffusion terms into a

high-probability safety constraint using a generator Upper Confidence Bound (UCB).
- We formulate a minimally invasive, safety-certified controller as a convex Quadratic Program (QP) that respects both aleatoric risk and epistemic uncertainty, supported by a formal inter-sample robustness margin.

## II. Preliminaries and Problem Formulation

### A. System Dynamics and Notation

We model the system by an Itô SDE:

$$\mathrm{d}\boldsymbol{x}_t = \mu(\boldsymbol{x}_t, \boldsymbol{u}_t)\,\mathrm{d}t + \Sigma(\boldsymbol{x}_t, \boldsymbol{u}_t)\,\mathrm{d}\boldsymbol{W}_t, \tag{1}$$

where the state $\boldsymbol{x}_t \in \mathbb{R}^n$, the control $\boldsymbol{u}_t \in \mathcal{U} \subseteq \mathbb{R}^m$, and $\boldsymbol{W}_t$ is a standard $p$-dimensional Wiener process on a complete probability space. $\mu$ and $\Sigma$ are the drift and diffusion terms, respectively. We assume $\mu$ and $\Sigma$ are locally Lipschitz. For the system operating within a compact set $\mathcal{D} \supset \mathcal{S}$, this ensures the existence of a strong solution. We often utilize the control-affine form:

$$\mathrm{d}\boldsymbol{x}_t = (f(\boldsymbol{x}_t) + g(\boldsymbol{x}_t)\boldsymbol{u}_t)\,\mathrm{d}t + \sigma(\boldsymbol{x}_t)\,\mathrm{d}\boldsymbol{W}_t. \tag{2}$$

We define the safe set $\mathcal{S}$ via a twice continuously differentiable barrier function $B \in C^2$:

$$\mathcal{S} := \{\boldsymbol{x} \in \mathbb{R}^n : 0 \le B(\boldsymbol{x}) < 1\}. \tag{3}$$

We assume the initial state $\boldsymbol{x}_0$ satisfies $B(\boldsymbol{x}_0) \le \gamma < 1$. For a digital implementation, we consider a finite time horizon $T$ discretized into $N$ steps of duration $\Delta t$, such that $T = N\Delta t$. The discrete sampling instants are denoted $t_k = k\Delta t$ for $k \in \{0, \dots, N\}$, and the sampled barrier value is $B_k := B(\boldsymbol{x}_{t_k})$.

### B. Stochastic Analysis and RA-CBFs

The generator $\mathcal{A}$ of $\boldsymbol{x}_t$ characterizes the expected rate of change of a function along the trajectories of (1). For $B \in C^2$ (the space of twice continuously differentiable functions), the generator's action on $B$ is denoted $\Gamma_B(\boldsymbol{x}, \boldsymbol{u}) := (\mathcal{A}B)(\boldsymbol{x}, \boldsymbol{u})$, and is given by:

$$\Gamma_B(\boldsymbol{x}, \boldsymbol{u}) = \mu(\boldsymbol{x}, \boldsymbol{u})^\top \nabla B(\boldsymbol{x}) + \tfrac{1}{2}\,\mathrm{Tr}\big(\Sigma\Sigma^\top \nabla^2 B(\boldsymbol{x})\big). \tag{4}$$

By Itô's formula, the evolution of the barrier function decomposes into a finite-variation (drift) term $I_L(t)$ and a stochastic (martingale) term $I_S(t)$, defined as:

$$I_L(t) := \int_0^t \Gamma_B(\boldsymbol{x}_s, \boldsymbol{u}_s)\,\mathrm{d}s, \quad I_S(t) := \int_0^t L_\Sigma B(\boldsymbol{x}_s, \boldsymbol{u}_s)\,\mathrm{d}W_s.$$

The barrier value at time $T$ is thus given by:

$$B(\boldsymbol{x}_T) = B(\boldsymbol{x}_0) + I_L(T) + I_S(T). \tag{5}$$

The term $L_\Sigma B(\boldsymbol{x}, \boldsymbol{u}) := \Sigma(\boldsymbol{x}, \boldsymbol{u})^\top \nabla B(\boldsymbol{x})$ represents the sensitivity of the barrier function to the stochastic noise.

The RA-CBF framework [14] bounds the probability that the stochastic term $I_S$ causes a barrier crossing. This requires bounding the noise sensitivity:

$$\|L_\Sigma B(\boldsymbol{x}, \boldsymbol{u})\|_2^2 \;\le\; \eta^2 \tag{6}$$

Under this bound, the boundary-crossing probability over horizon $T$ from initial condition $\gamma$ is bounded via a DDS time-change and the reflection principle (see, e.g., [21], [22]):

$$\mathbb{P}\big(\max_{t \le T} B(\boldsymbol{x}_t) \ge 1\big) \le \mathrm{erfc}\left(\frac{1-\gamma}{\sqrt{2}\,\eta\,\sqrt{T}}\right). \tag{7}$$

Inverting for a target risk threshold $\rho_d$ yields the required continuous-time stochastic safety margin:

$$\Delta_{\mathrm{CT}}^\star(\rho_d, \eta, T) = \sqrt{2}\,\eta\,\sqrt{T}\;\mathrm{erfc}^{-1}(\rho_d). \tag{8}$$

Here, $\Delta_{\mathrm{CT}}^\star$ is the minimum buffer that must be reserved for the stochastic term over $[0, T]$. Intuitively, the DDS time change rewrites the martingale term as a standard Brownian motion evaluated at its quadratic variation, and the reflection principle converts the resulting boundary-crossing event into the Gaussian tail in (7). The continuous-time RA-CBF design then enforces that the deterministic accumulation $I_L$ leaves this margin available:

$$I_L \;\le\; (1-\gamma) - \Delta_{\mathrm{CT}}^\star(\rho_d, \eta, T). \tag{9}$$

### C. Problem Formulation

We consider the safety-critical control of the stochastic system (1) under a digital controller with sample period $\Delta t$ over the finite horizon $T$. The dynamics $f$ and $\sigma$ (or $D = \sigma\sigma^\top$) are unknown and learned via Bayesian methods. We denote the sampled values of the integrated drift and stochastic terms as $I_{L,k} := I_L(t_k)$ and $I_{S,k} := I_S(t_k)$, respectively.

**Problem II.1.** Given a safe set $\mathcal{S}$ defined by (3), an initial state $\boldsymbol{x}_0 \in \mathcal{S}$, a horizon $T$, and a total risk requirement $\rho_{\mathrm{req}} \in (0, 1)$. Design a discrete-time feedback controller $\boldsymbol{u}_k = \kappa(\boldsymbol{x}_k)$ that enforces the RA-CBF constraint, accounting for discretization effects and uncertainty in the learned dynamics, such that the probability of failure $\mathbb{P}(\exists\, k \in \{0, \dots, N\} : \boldsymbol{x}_k \notin \mathcal{S})$ is bounded by $\rho_{\mathrm{req}}$ with high confidence.

To solve this problem, the total risk requirement $\rho_{\mathrm{req}}$ is allocated to different sources of uncertainty. We designate $\rho_d$ as the internal risk threshold for the system's inherent stochasticity (aleatoric uncertainty), which is used to calculate the continuous-time and discrete-time safety margins introduced below. Additional thresholds, such as $\delta_H$ and $\delta_\epsilon$, are reserved for model uncertainty (epistemic) and inter-sample effects. The formal mapping between these thresholds and the total requirement $\rho_{\mathrm{req}}$ is given in Thm. V.5 and Prop. V.6.

**Example II.2.** As a running example, consider a mobile robot that must remain inside a circular constraint region of radius $R_c > 0$ centered at the origin $s_0$ of an inertial frame $\mathcal{I}$. We choose $R_c = 1$. For a nominal "task drive," we use a simple outward radial policy; the safety filter ensures the robot remains within the constraint set. The robot is modeled as a 2D stochastic single-integrator,

$$\mathrm{d}\boldsymbol{x}_t = \boldsymbol{u}_t\,\mathrm{d}t + \Sigma\,\mathrm{d}\boldsymbol{W}_t, \tag{10}$$
$$\boldsymbol{x}_t = \begin{bmatrix} x \\ y \end{bmatrix} \in \mathbb{R}^2, \quad \boldsymbol{u}_t = \begin{bmatrix} v_x \\ v_y \end{bmatrix} \in \mathbb{R}^2,$$

with $\boldsymbol{W}_t$ a 2D Wiener process and $\Sigma = \mathrm{diag}(\sigma_x, \sigma_y)$. In Fig. 1, we use isotropic diffusion $\sigma_x = \sigma_y = \sigma = 0.1$.

We take the circular barrier $B(\boldsymbol{x}) = \frac{\|\boldsymbol{x}\|^2}{R_c^2}$ (cf. (3)). Then $\nabla B(\boldsymbol{x}) = \frac{2}{R_c^2}\boldsymbol{x}$, $\nabla^2 B(\boldsymbol{x}) = \frac{2}{R_c^2} I$, and the generator (4) is

$$\begin{aligned}\Gamma_B(\boldsymbol{x}, \boldsymbol{u}) &= \nabla B(\boldsymbol{x})^\top \boldsymbol{u} + \tfrac{1}{2}\,\mathrm{Tr}\big(\Sigma\Sigma^\top \nabla^2 B(\boldsymbol{x})\big) \\ &= \nabla B(\boldsymbol{x})^\top \boldsymbol{u} + \frac{\sigma_x^2 + \sigma_y^2}{R_c^2},\end{aligned}$$

which reduces to $\Gamma_B(\boldsymbol{x}, \boldsymbol{u}) = \nabla B(\boldsymbol{x})^\top \boldsymbol{u} + \frac{2\sigma^2}{R_c^2}$ for the isotropic case. For the noise sensitivity bound (6) with isotropic diffusion ($\Sigma = \sigma I$), we have $\|\nabla B(\boldsymbol{x})\|_2 = \frac{2}{R_c^2}\|\boldsymbol{x}\|_2 \le \frac{2}{R_c}$ on the safe set $\{\boldsymbol{x} : B(\boldsymbol{x}) < 1\}$, so $\|\Sigma^\top \nabla B(\boldsymbol{x})\|_2 \le \frac{2\sigma}{R_c}$.

## III. Discrete-Time RA-CBF via Martingale Concentration

The continuous-time RA-CBF framework provides a foundation for safety, but physical controllers operate in discrete time. We now extend this analysis to sampled-data systems, addressing safety at discrete instants $t_k = k\Delta t$ w.r.t. $(\mathcal{F}^k)_{k\ge 0}$ (representing the history of the process up to time $t_k$). To this end, we derive two distinct safety guarantees tailored to different assumptions on the system's stochastic behavior. First, we establish a robust, time-uniform guarantee using general martingale concentration bounds, which relies on a predictable quadratic variation threshold over the horizon. Second, we derive a tighter bound by analyzing the underlying continuous-time stochastic integral, a pathwise approach that can recover the non-conservative performance of the continuous-time design under a stronger noise sensitivity assumption.

Let $(\mathcal{F}_t)_{t\ge 0}$ denote the completed Brownian filtration and write $\mathcal{F}^i := \mathcal{F}_{t_i}$. To analyze the system at discrete times $t_k$, we consider the one-step increment of the martingale term, defined as $\Delta I_{S,i} := \int_{t_i}^{t_{i+1}} L_\Sigma B(\boldsymbol{x}_t, \boldsymbol{u}_t)\,\mathrm{d}\boldsymbol{W}_t$. By the Itô isometry, its conditional variance (the predictable quadratic variation increment) is:

$$V_i := \mathbb{E}\big[(\Delta I_{S,i})^2 \mid \mathcal{F}^i\big] = \mathbb{E}\left[\int_{t_i}^{t_{i+1}} \|L_\Sigma B(\boldsymbol{x}_t, \boldsymbol{u}_t)\|_2^2\,\mathrm{d}t \;\middle|\; \mathcal{F}^i\right].$$

The total predictable quadratic variation is the sum of these single-step expected variances, $V_N = \sum V_i$. To establish the necessary concentration bounds, we rely on the continuous-time pathwise noise sensitivity bound $\|L_\Sigma B\|_2^2 \le \eta^2$ from (6). This pathwise bound ensures that both the predictable variation and the realized quadratic variation (the actual accumulated variance) are bounded by the total accumulated threshold:

$$V_N \le \eta^2 T, \quad \text{and} \quad \int_0^T \|L_\Sigma B(\boldsymbol{x}_t, \boldsymbol{u}_t)\|_2^2\,\mathrm{d}t \le \eta^2 T\,. \tag{11}$$

The available deterministic buffer is $a := 1 - \gamma - \sup_{k\le N} I_{L,k}$.

**Lemma III.1** (Exponential supermartingales for $I_{S,k}$)**.** Define $\Delta I_{S,i} := \int_{t_i}^{t_{i+1}} L_\Sigma B(\boldsymbol{x}_s, \boldsymbol{u}_s)\,dW_s$, $S_k := \sum_{j=0}^{k-1} \Delta I_{S,j}$, and $\Delta[I_S]_i := \int_{t_i}^{t_{i+1}} \|L_\Sigma B(\boldsymbol{x}_s, \boldsymbol{u}_s)\|_2^2\,ds$.

*(a)* Realized Quadratic Variation form. The process

$$\widehat{M}_k(\lambda) := \exp\left\{\lambda S_k - \tfrac{\lambda^2}{2}\sum_{j=0}^{k-1} \Delta[I_S]_j\right\}$$

is a nonnegative supermartingale w.r.t. $(\mathcal{F}^k)_{k\ge 0}$ for any $\lambda \in \mathbb{R}$.

*(b) Pathwise deterministic proxy.* If $\|L_\Sigma B(\boldsymbol{x}_t, \boldsymbol{u}_t)\|_2 \le \eta$ a.s. for all $t \in [0, T]$, then

$$M_k(\lambda) := \exp\left\{\lambda S_k - \tfrac{\lambda^2}{2}\eta^2 t_k\right\}$$

is a nonnegative supermartingale.

*Proof.* For (a), $\mathcal{E}_t(\lambda) := \exp\{\lambda I_S(t) - \frac{1}{2}\lambda^2 [I_S]_t\}$ is the Doléans–Dade exponential of a continuous local martingale, hence a nonnegative supermartingale; sampling at $t_k$ gives $\widehat{M}_k$ [23, Prop. 5.21 and p. 136]. For (b), we confirm the supermartingale property by analyzing the one-step conditional expectation. Factoring out the known terms yields $\mathbb{E}[M_{k+1}(\lambda) \mid \mathcal{F}^k] = M_k(\lambda) \cdot \mathbb{E}[e^{\lambda\Delta I_{S,k}} \mid \mathcal{F}^k] e^{-\frac{\lambda^2}{2}\eta^2\Delta t}$. The pathwise bound on $L_\Sigma B$ ensures the increment $\Delta I_{S,k}$ is conditionally sub-Gaussian, satisfying the inequality $\mathbb{E}[e^{\lambda\Delta I_{S,k}} \mid \mathcal{F}^k] \le e^{\frac{\lambda^2}{2}\eta^2\Delta t}$. Substituting this bound into the expression above proves the condition $\mathbb{E}[M_{k+1}(\lambda) \mid \mathcal{F}^k] \le M_k(\lambda)$. □

With the supermartingale property established by Lemma III.1, applying Ville's inequality [24] yields the following safety guarantee:

**Proposition III.2** (Discrete-time RA-CBF)**.** If $a > 0$ and either (i) the pathwise bound $\|L_\Sigma B\|_2 \le \eta$ holds, or (ii) the realized quadratic variation satisfies $\sum_{i=0}^{N-1} \Delta[I_S]_i \le \eta^2 T$, then

$$\mathbb{P}\left(\max_{0\le k\le N} B_k \ge 1\right) \le \exp\left(-\frac{a^2}{2\,\eta^2\,T}\right). \tag{12}$$

*Proof.* A barrier violation occurs if $B_k \ge 1$ for some $k \in \{0, \dots, N\}$. By decomposing the barrier value as $B_k = \gamma + I_{L,k} + I_{S,k}$, a violation implies:

$$\gamma + I_{L,k} + I_{S,k} \ge 1 \implies I_{S,k} \ge 1 - \gamma - I_{L,k}.$$

Let $a := 1 - \gamma - \sup_{0\le k\le N} I_{L,k}$ be the minimum available deterministic buffer over the horizon. A barrier violation implies the stochastic part must have exceeded this buffer, which establishes the inclusion of events:

$$\left\{\max_{0\le k\le N} B_k \ge 1\right\} \subseteq \left\{\max_{0\le k\le N} I_{S,k} \ge a\right\}.$$

Consequently, $\mathbb{P}(\max_k B_k \ge 1) \le \mathbb{P}(\max_k I_{S,k} \ge a)$. Let $S_k := I_{S,k}$. We now show that both conditions (i) and (ii) yield the same bound on the probability of the event $\{\max_k S_k \ge a\}$.

*Case (i): Pathwise bound.* Under this condition, we use Lemma III.1(b), which states that $M_k(\lambda) := \exp\{\lambda S_k - \frac{\lambda^2}{2}\eta^2 t_k\}$ is a nonnegative supermartingale. Applying Ville's inequality directly yields, for any $\lambda > 0$,

$$\mathbb{P}\left(\max_{0\le k\le N} S_k \ge a\right) \le \exp\left(-\lambda a + \tfrac{\lambda^2}{2}\eta^2 T\right).$$

The right-hand side is minimized by choosing $\lambda^\star = a/(\eta^2 T)$, giving the bound in (12).

*Case (ii): Realized quadratic variation.* Under this condition, we use Lemma III.1(a), where $\widehat{M}_k(\lambda) := \exp\{\lambda S_k - \frac{\lambda^2}{2}\sum_{j=0}^{k-1}\Delta[I_S]_j\}$ is a nonnegative supermartingale with $\mathbb{E}[\widehat{M}_0] = 1$. The condition (ii) states that $\sum_{j=0}^{k-1}\Delta[I_S]_j \le \eta^2 T$. This establishes the event inclusion:

$$\left\{\max_{0\le k\le N} S_k \ge a\right\} \subseteq \left\{\max_{0\le k\le N} \widehat{M}_k(\lambda) \ge e^{\lambda a - \frac{\lambda^2}{2}\eta^2 T}\right\}.$$

Applying Ville's inequality to the supermartingale $\widehat{M}_k(\lambda)$ gives, for any $\lambda > 0$,

$$\mathbb{P}\left(\max_{0\le k\le N} S_k \ge a\right) \le \frac{1}{e^{\lambda a - \frac{\lambda^2}{2}\eta^2 T}} = \exp\left(-\lambda a + \tfrac{\lambda^2}{2}\eta^2 T\right).$$

Optimizing this bound at $\lambda^\star = a/(\eta^2 T)$ again yields the result $\exp(-a^2/(2\eta^2 T))$. □

Setting the right-hand side of (12) to be $\le \rho_d$ and solving for the required buffer yields the discrete-time stochastic safety margin

$$\Delta^\star_{\mathrm{DT}} := \sqrt{2\,\eta^2 T\,\ln(1/\rho_d)}, \qquad I_{L,N} \le (1-\gamma) - \Delta^\star_{\mathrm{DT}}. \quad (13)$$

*Remark* III.3 (Guarantee and Margin Comparison). This approach provides a time-uniform, pathwise guarantee on a single system trajectory by conditioning on the process history, distinct from methods that only bound moments of the state distribution [14]. The resulting discrete-time margin, $\Delta^\star_{\mathrm{DT}}$, scales with $\sqrt{\ln(1/\rho_d)}$ and is generally more conservative than its continuous-time counterpart, $\Delta^\star_{\mathrm{CT}}$, which scales with $\mathrm{erfc}^{-1}(\rho_d)$. As shown in Prop. IV.1, this conservatism is removed under a stronger pathwise assumption, recovering the tighter CT performance.

## IV. Tighter Discrete-Time Bounds Under Pathwise Noise Sensitivity Assumption

While robust, the guarantee from Proposition III.2 can be conservative because it relies on general-purpose inequalities. We obtain a tighter bound by exploiting that the samples $\{B_k\}$ are snapshots of a continuous-time Itô SDE. Analyzing the underlying stochastic integral directly adapts the continuous-time RA-CBF framework to discrete sampling and recovers the continuous-time performance under a stronger, pathwise assumption on noise sensitivity.

**Proposition IV.1** (Tighter Discrete-Time RA-CBF Bound)**.** Assume the pathwise noise sensitivity bound $\|L_\Sigma B\|_2 \le \eta$ holds for all $t \in [0,T]$. Let the finite-variation part of the barrier evolution be bounded by a deterministic constant $b$ such that $\sup_{k\le N} I_{L,k} \le b$. Define the available deterministic buffer as $a := 1-\gamma-b$. If $a > 0$, then the probability of a barrier violation at any discrete sampling instant is bounded by:

$$\mathbb{P}\left(\max_{0\le k\le N} B_k \ge 1\right) \le \mathrm{erfc}\left(\frac{a}{\eta\sqrt{2T}}\right). \quad (14)$$

*Proof.* The proof relates the discrete crossing event to the maximum of the continuous-time martingale, which is then bounded by a time-change argument and the reflection principle for Brownian motion.

First, a barrier crossing at a discrete step $k$ implies $I_{S,k} \ge 1-\gamma-I_{L,k} \ge 1-\gamma-b = a$. The maximum over discrete samples is bounded by the continuous-time supremum, giving the event inclusion:

$$\left\{\max_{k\le N} B_k \ge 1\right\} \subseteq \left\{\sup_{t\in[0,T]} I_S(t) \ge a\right\},$$

and thus $\mathbb{P}(\max_{k\le N} B_k \ge 1) \le \mathbb{P}(\sup_{t\in[0,T]} I_S(t) \ge a)$.

Next, we bound this supremum. The process $I_S(t)$ is a continuous local martingale. Its quadratic variation is $[I_S]_t = \int_0^t \|L_\Sigma B(\boldsymbol{x}_s,\boldsymbol{u}_s)\|_2^2 \mathrm{d}s$. The pathwise assumption $\|L_\Sigma B\|_2 \le \eta$ ensures this is bounded:

$$[I_S]_T = \int_0^T \|L_\Sigma B(\boldsymbol{x}_s,\boldsymbol{u}_s)\|_2^2 \mathrm{d}s \le \eta^2 T$$

By the Dambis-Dubins-Schwarz (DDS) Theorem [21], $I_S(t)$ can be represented as a standard one-dimensional Brownian motion $\mathcal{W}$ evaluated at the time given by its quadratic variation: $I_S(t) = \mathcal{W}([I_S]_t)$.

This representation allows us to relate the supremum of $I_S(t)$ to the supremum of $\mathcal{W}(\tau)$. The equality is in distribution $(d)$, and the subsequent inequality is one of stochastic dominance $(st)$, since $[I_S]_T$ is a random variable bounded by the constant $\eta^2 T$ and the supremum of a Brownian motion is non-decreasing with its time horizon:

$$\sup_{t\in[0,T]} I_S(t) \overset{d}{=} \sup_{0\le\tau\le[I_S]_T} \mathcal{W}(\tau) \le_{st} \sup_{0\le\tau\le\eta^2 T} \mathcal{W}(\tau).$$

Finally, we apply the reflection principle to the stochastically dominant term. For a standard Brownian motion $\mathcal{W}(\tau)$ and a level $a > 0$, the probability of the supremum exceeding this level is given by the reflection principle [21, Prop 3.7]:

$$\mathbb{P}\left(\sup_{0\le\tau\le S} \mathcal{W}(\tau) \ge a\right) = 2\mathbb{P}(\mathcal{W}(S) \ge a) = \mathrm{erfc}\left(\frac{a}{\sqrt{2S}}\right).$$

Setting the time duration $S = \eta^2 T$ and combining all steps yields the final bound. □

*Remark* IV.2 (Equivalence to Continuous-Time Margin). Inverting the bound in (14) for a risk threshold $\rho_d$ provides the required stochastic safety margin $a = \eta\sqrt{2T}\cdot\mathrm{erfc}^{-1}(\rho_d)$. This is identical to the continuous-time margin $\Delta^\star_{\mathrm{CT}}$ from Eq. (8). This confirms that, under the stronger pathwise assumption, a discrete-time controller can be designed with the same non-conservative margin as a continuous-time one, avoiding the pessimism of general martingale bounds.

### A. Illustrative Margin Comparison

To make the gap between the two margins concrete before turning to implementation, consider a system with initial state $\gamma = 0.1$, horizon $T = 10$s, noise sensitivity $\eta = 0.08$ units/$\sqrt{\text{s}}$, and risk threshold $\rho_d = 0.10$. The original discrete margin, derived from a general martingale bound (Prop. III.2), is

$\Delta^\star_{\mathrm{DT}} = \sqrt{2\eta^2 T \ln(1/\rho_d)} = \sqrt{2(0.08^2)(10)\ln(10)} \approx 0.5429$. This yields an allowable drift of $I_{L,N} \le (1-0.1) - 0.5429 = 0.3571$. In contrast, the tighter discrete margin (Prop. IV.1), obtained via a pathwise approach that recovers the continuous-time margin $\Delta^\star_{\mathrm{CT}}$, is $\Delta^\star_{\mathrm{CT}} = \sqrt{2}\,\eta\,\sqrt{T}\ \mathrm{erfc}^{-1}(\rho_d) \approx 0.4161$. The corresponding allowable drift increases to $I_{L,N} \le (1-0.1) - 0.4161 = 0.4839$. Applying this tighter bound to the system from Example II.2 (see Fig. 1) provides a significant reduction in conservatism. The required stochastic safety margin is reduced by over 23%, which directly translates into a 35% larger threshold for the deterministic drift. This enables a controller that is less invasive while providing the same probabilistic safety guarantee.

*Remark* IV.3 (Tighter Bounds due to stronger Assumptions). This performance gain comes from an important trade-off in the underlying assumptions.

- The bound in Prop. III.2 only requires that the *total expected variance* over the horizon is bounded: $\sum \mathbb{E}[(\Delta I_{S,i})^2 \mid \mathcal{F}^i] \le \eta^2 T$. This allows for periods where the instantaneous noise sensitivity $\|L_\Sigma B\|$ might exceed $\eta$, as long as it is compensated for at other times.
- The tighter bound in Prop. IV.1 requires a much stricter, *pathwise uniform bound*: $\|L_\Sigma B(\boldsymbol{x}_t, \boldsymbol{u}_t)\|_2 \le \eta$ for all $t \in [0, T]$.

However, Propositions III.2 and IV.1 ensure safety only at sampling instants $t_k$. Under a Zero-Order Hold (ZOH), the state evolves between samples, so inter-sample violations are possible. We address this problem in the next section.

## V. Bayesian RA-CBF Implementation

### A. Receding Horizon Control with Learned Dynamics

We use Receding Horizon Control (RHC) for a practical, real-time implementation. At each step, the controller enforces a safety constraint over a finite lookahead horizon. This requires propagating the posterior uncertainty of the learned dynamics through the generator.

*a) Generator posterior:* For a barrier function $B \in C^2$ with gradient $\nabla B(x)$ and Hessian $H_B(x) = \nabla^2 B(x)$, the generator is

$$\Gamma_B(x,u) := \nabla B(x)^\top f(x) + L_g B(x)\, u + \tfrac{1}{2}\ \mathrm{Tr}\big(H_B(x) D(x)\big),$$

where $L_g B(x) := \nabla B(x)^\top g(x)$. We approximate the posteriors of $f$ and $\mathrm{vec}(D)$ by independent Gaussians. Since the term $\ell(D) := \frac{1}{2}\mathrm{Tr}(H_B D) = \frac{1}{2}\mathrm{vec}(H_B)^\top \mathrm{vec}(D)$ is a linear functional of $D$, the generator posterior is also approximately Gaussian:

$$\Gamma_B(x,u) \mid \mathcal{D}_t \approx \mathcal{N}\big(\mu_\Gamma(x,u), \sigma^2_\Gamma(x)\big). \tag{15}$$

The moments are calculated as follows:

$$\mu_\Gamma(x,u) = \nabla B(x)^\top \mu_f(x) + L_g B(x)\, u + \tfrac{1}{2}\ \mathrm{Tr}\big(H_B(x)\, \mu_D(x)\big), \tag{16a}$$

$$\sigma^2_\Gamma(x) = \nabla B(x)^\top \Sigma_f(x)\, \nabla B(x) + \tfrac{1}{4}\ \big(\mathrm{vec}\, H_B(x)\big)^\top \Sigma_D(x)\ \big(\mathrm{vec}\, H_B(x)\big). \tag{16b}$$

The true diffusion matrix must be positive semi-definite, $D(x) \succeq 0$. By using an *unconstrained* Gaussian posterior for $\mathrm{vec}(D)$, we are considering a superset of possible diffusion matrices. A UCB for the linear functional $\ell(D)$ computed with this posterior is therefore provably conservative (i.e., larger) than one computed under the strict $D \succeq 0$ constraint.

*b) UCB for the generator:* For a confidence level $1-\delta_H$, the generator's Upper Confidence Bound (UCB) is constructed using the Gaussian quantile $q_{\delta_H} = \Phi^{-1}(1-\delta_H)$, where $\Phi$ is the standard normal CDF:

$$\mathrm{UCB}_\Gamma(x,u;\delta_H) := \mu_\Gamma(x,u) + q_{\delta_H}\, \sigma_\Gamma(x). \tag{17}$$

To keep the controller formulas agnostic to the certificate used, define the stochastic safety margin

$$\Delta_\rho := \begin{cases} \Delta^\star_{\mathrm{DT}}, & \text{when Prop. III.2 is used,} \\ \Delta^\star_{\mathrm{CT}}, & \text{when Prop. IV.1 is used.} \end{cases}$$

**Definition V.1** (RHC Bayesian RA-CBF)**.** Let $T = N\Delta t$ be the lookahead horizon, and let $\delta_H$ and $\delta_\epsilon$ be the confidence thresholds for model uncertainty and inter-sample robustness, respectively. A control sequence $\{\boldsymbol{u}_j\}$ is a Bayesian RA-CBF policy starting at $t_k$ if, for $j \in \{k, \dots, k+N-1\}$, $\boldsymbol{u}_j \in \mathcal{U}$ satisfies:

$$\mathrm{UCB}_\Gamma(\boldsymbol{x}_j, \boldsymbol{u}_j; \delta_H/N) + \epsilon(\Delta t, \delta_\epsilon/N) \le \frac{1 - B(\boldsymbol{x}_k) - \Delta_\rho}{T}, \tag{18}$$

where the available safety buffer is $1 - B(\boldsymbol{x}_k)$. The per-step confidence $1-\delta_H/N$ ensures the UCB holds jointly over the horizon with probability at least $1-\delta_H$ via a union bound.

The objective is to derive a high-probability bound, $\epsilon(\Delta t, \delta'_\epsilon)$, on the maximum deviation of the generator's Upper Confidence Bound (UCB) over a single sampling interval $[t_k, t_{k+1})$. Under a Zero-Order Hold (ZOH), the control input $\boldsymbol{u}_k$ is constant during this time. Specifically, the goal is to find $\epsilon$ such that for a single-step risk threshold $\delta'_\epsilon := \delta_\epsilon/N \in (0,1)$, the following holds:

$$\Pr\left( \sup_{t \in [t_k, t_{k+1})} |UCB_\Gamma(\boldsymbol{x}_t, \boldsymbol{u}_k) - UCB_\Gamma(\boldsymbol{x}_k, \boldsymbol{u}_k)| \ge \epsilon(\Delta t, \delta'_\epsilon) \right) \le \delta'_\epsilon.$$

*1) Assumptions:* The derivation relies on two standard assumptions holding over a compact set $\mathcal{K} \subset \mathcal{S}$ where the system operates.

*Assumption* V.2 (Lipschitz UCB). The UCB function used in the controller, $UCB_\Gamma(x,u) = \mu_\Gamma(x,u) + q_{\delta_H/N}\sigma_\Gamma(x)$, is locally Lipschitz continuous in state $x$ with constant $L_U$. A continuously differentiable kernel, such as the Squared Exponential, produces a GP posterior whose mean and standard deviation functions are also continuously differentiable. Any continuously differentiable function is locally Lipschitz, thereby satisfying the assumption [25, Ch. 2.2–2.3].

*Assumption* V.3 (Bounded dynamics). The system's drift $f(x,u)$ and diffusion $\sigma(x,u)$ are uniformly bounded by constants $M_\mu := \sup \|f(x,u)\|_2$ and $M_\Sigma := \sup \|\sigma(x,u)\|_2$, where $\|\cdot\|_2$ denotes the Euclidean norm for vectors and the spectral (operator) norm for matrices.

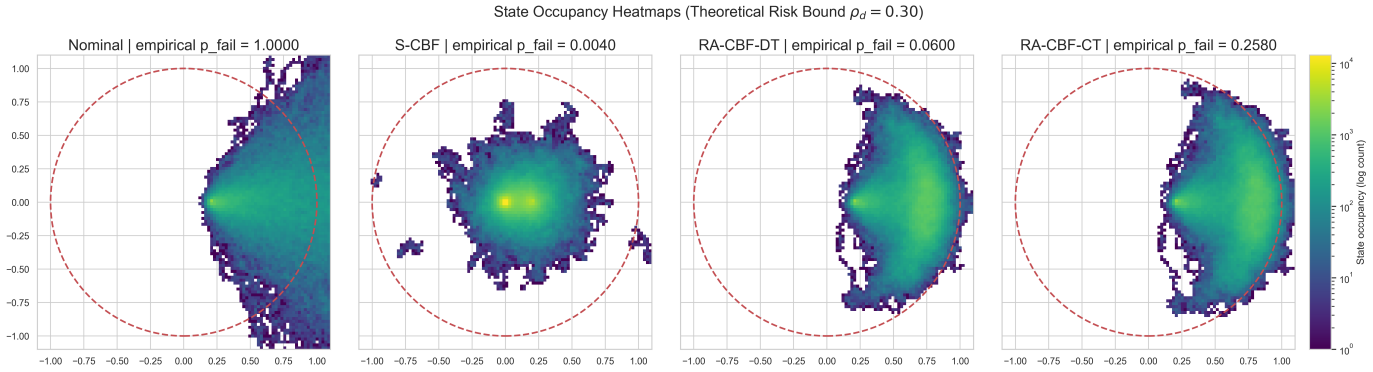


Fig. 1. State occupancy heatmaps (log-count) for Nominal, S-CBF, RA-CBF-DT, and RA-CBF-CT at $v_{\max} = 0.4$. Each panel shows empirical state visitation density across 1000 rollouts; the dashed red circle denotes the safety boundary $B(x) = 1$ (radius $R_c$). Panel titles report the empirical probability of failure $\hat{p}_{\text{fail}}$ and the design risk bound $\rho_d$ (cf. Prop. III.2, Prop. IV.1, and Eqs. (8)–(13)). RA-CBF variants concentrate mass closer to the boundary relative to Nominal and S-CBF; the discrete-time variant tends to be more conservative. The tighter the match between empirical and theoretical bounds, the greater the utility for risk-aware control. For one execution of this system, see video.

We start by relating the UCB deviation to state deviation. The Lipschitz property of the UCB provides the starting point. It links the deviation of the UCB to the deviation of the state:

$$|UCB_\Gamma(\boldsymbol{x}_t, \boldsymbol{u}_k) - UCB_\Gamma(\boldsymbol{x}_k, \boldsymbol{u}_k)| \leq L_U \|\boldsymbol{x}_t - \boldsymbol{x}_k\|. \tag{19}$$

This inequality shows that if we can bound how far the state $\boldsymbol{x}_t$ can wander from its starting point $\boldsymbol{x}_k$, we can bound the UCB's deviation. Next, we decompose the state deviation. The state's evolution is governed by the integral form of the SDE. Using the triangle inequality, we can separate the deviation into its deterministic (drift) and stochastic (diffusion) parts:

$$\|\boldsymbol{x}_t - \boldsymbol{x}_k\| \leq \left\| \int_{t_k}^{t} f(\boldsymbol{x}_s, \boldsymbol{u}_k)\, ds \right\| + \left\| \int_{t_k}^{t} \sigma(\boldsymbol{x}_s, \boldsymbol{u}_k)\, dW_s \right\|. \tag{20}$$

Next, we bound the deterministic drift term using the bounded dynamics assumption (Assumption V.3). The integral is bounded by the maximum drift rate $M_\mu$ multiplied by the elapsed time:

$$\left\| \int_{t_k}^{t} f(\boldsymbol{x}_s, \boldsymbol{u}_k) ds \right\| \leq \int_{t_k}^{t} \|f(\boldsymbol{x}_s, \boldsymbol{u}_k)\| ds \leq M_\mu (t - t_k) \leq M_\mu \Delta t.$$

The diffusion term, $M_t := \int_{t_k}^{t} \sigma(\boldsymbol{x}_s, \boldsymbol{u}_k) dW_s$, is a vector-valued martingale. Bounding the norm of a vector martingale directly is difficult. We simplify the problem by relating the vector's norm to its maximum projection onto a finite set of direction vectors, known as a covering net. For a $1/2$-net $N$ of the unit sphere $\mathbb{S}^{n-1}$, any vector $z \in \mathbb{R}^n$ satisfies $\|z\| \leq 2 \max_{v \in N} v^\top z$ [26, Ex. 4.4.2]. Applying this to $M_t$ means that if $\|M_t\| \geq r$, then at least one scalar projection $v^\top M_t$ must be greater than or equal to $r/2$. Next, we bound the tail of each scalar projection $m_t^v := v^\top M_t$, which is a continuous local martingale. By the Dambis-Dubins-Schwarz theorem, we can represent $m_t^v$ as a standard Brownian motion $B$ run on a different "clock," where the new time $\tau$ is the quadratic variation of $m_t^v$ [21, Thm. V.1.6]. The total duration of this clock is bounded: $\tau_{max} \leq M_\Sigma^2 \Delta t$. Using the *reflection principle* for Brownian motion, the probability that the supremum of $B$ over a time $S$ exceeds a level $a$ is given by $\operatorname{erfc}(a/\sqrt{2S})$ [22, §2.2]. This gives us a tail probability for a single scalar projection exceeding a level $a$:

$$\mathbb{P}\left( \sup_{t \in [t_k, t_{k+1})} m_t^v \geq a \right) \leq \operatorname{erfc}\left( \frac{a}{\sqrt{2M_\Sigma^2 \Delta t}} \right).$$

Now, we combine the bounds for all scalar projections using a *union bound*. The number of vectors in the net is bounded by $|N| \leq 5^n$ [26, Cor. 4.2.13].

$$\mathbb{P}(\sup_t \|M_t\| \geq r) \leq \sum_{v \in N} \mathbb{P}(\sup_t m_t^v \geq r/2) \leq 5^n \cdot \operatorname{erfc}\left( \frac{r/2}{\sqrt{2M_\Sigma^2 \Delta t}} \right).$$

We set this probability less than or equal to our risk threshold $\delta'_\epsilon$. To solve for $r$, we use the Gaussian tail bound $\operatorname{erfc}(z) \leq 2e^{-z^2}$. Solving the resulting inequality for $r$ gives the high-probability bound on the diffusion term:

$$r(\Delta t, \delta'_\epsilon) := 2\sqrt{2M_\Sigma^2 \Delta t \left( \ln\left( \frac{2}{\delta'_\epsilon} \right) + n \ln 5 \right)}.$$

This bound holds with probability at least $1 - \delta'_\epsilon$. The term $n \ln 5$ arises from the size of the covering net and captures the dependence on the state dimension $n$.

Finally, we combine the deterministic drift bound and the high-probability diffusion bound to bound the total state deviation. We then substitute this into the Lipschitz inequality from Eq. (19).

With probability at least $1 - \delta'_\epsilon$:

$$\sup_{t \in [t_k, t_{k+1})} \|\boldsymbol{x}_t - \boldsymbol{x}_k\| \leq M_\mu \Delta t + r(\Delta t, \delta'_\epsilon).$$

Multiplying by the Lipschitz constant $L_U$ yields the final expression for the inter-sample margin.

**Proposition V.4** (Inter-Sample Margin)**.** Under the Lipschitz UCB and Bounded Dynamics assumptions, for any single-step risk threshold $\delta'_\epsilon \in (0, 1)$, the deviation of the UCB over the

interval $[t_k, t_{k+1})$ is bounded by $\epsilon(\Delta t, \delta'_\epsilon)$ with probability at least $1 - \delta'_\epsilon$, where:

$$\epsilon(\Delta t, \delta'_\epsilon) = L_U \left( M_\mu \Delta t + 2\sqrt{2M_\Sigma^2 \Delta t \left(\ln(2/\delta'_\epsilon) + n \ln 5\right)} \right).$$

To guarantee safety over a full horizon of $N$ steps with total risk $\delta_\epsilon$, we apply a union bound. We allocate a per-step risk threshold of $\delta'_\epsilon = \delta_\epsilon / N$. The probability that the margin fails in *any* of the $N$ steps is at most $N \cdot (\delta_\epsilon / N) = \delta_\epsilon$. Therefore, $\epsilon(\Delta t, \delta_\epsilon / N)$ is used in the receding-horizon controller to ensure safety with confidence $1 - \delta_\epsilon$ over the entire window. □

This RHC strategy provides a practical, step-by-step implementation for the controller described in Problem II.1.

### B. High-Confidence Probabilistic Safety Guarantee

By enforcing the condition in Def. V.1, we achieve a rolling probabilistic safety guarantee (see Fig. 2 for an illustration of the adaptive safety buffer).

**Theorem V.5** (RHC Probabilistic Safety Guarantee)**.** Assume a Bayesian RA-CBF strategy satisfying Def. V.1 is implemented at $t_k$. Let the risk thresholds be $\rho_d$ (stochastic), $\delta_H$ (model), and $\delta_\epsilon$ (inter-sample). If the generator UCB is locally Lipschitz in $\boldsymbol{x}$, the probability of a safety violation within the horizon is bounded:

$$\mathbb{P}\left( \sup_{t \in [t_k, t_k+T]} B(\boldsymbol{x}_t) \geq 1 \,\middle|\, \mathcal{D}_{t_k}, \boldsymbol{x}_k \right) \leq \rho_d + \delta_H + \delta_\epsilon - \rho_d(\delta_H + \delta_\epsilon). \tag{21}$$

*Proof.* We analyze the probability of failure $F := \{\sup_{t \in [t_k, t_k+T]} B(\boldsymbol{x}_t) \geq 1\}$. We recall the two key events introduced above.

Model Uncertainty Event ($E_U$): Let $E_{U,j}$ be the event that the true generator is bounded by the UCB at step $j$. Given the quantile $q_{\delta_H/N}$ used in Def. V.1, we have $\mathbb{P}(E_{U,j}^C) \leq \delta_H / N$. We define $E_U$ as the joint event that the UCB holds for all steps in the horizon: $E_U := \bigcap_{j=k}^{k+N-1} E_{U,j}$. By the union bound (Boole's inequality), $\mathbb{P}(E_U^C) = \mathbb{P}\left(\bigcup_{j=k}^{k+N-1} E_{U,j}^C\right) \leq \sum_{j=k}^{k+N-1} \mathbb{P}(E_{U,j}^C) \leq N \cdot \frac{\delta_H}{N} = \delta_H$.

Inter-Sample Robustness Event ($E_\epsilon$): Let $E_\epsilon$ be the joint event that the inter-sample margin $\epsilon(\Delta t, \delta_\epsilon / N)$ successfully bounds the deviation of the UCB (as defined in Assumption V.2) during all intervals $[t_j, t_{j+1})$. As established via a union bound in the derivation of the Inter-Sample Margin, $\mathbb{P}(E_\epsilon^C) \leq \delta_\epsilon$.

Combined Event ($E_G$): Let $E_G = E_U \cap E_\epsilon$. This is the event where both the model UCB and the inter-sample margin hold throughout the horizon. By union bound, $\mathbb{P}(E_G^C) \leq \mathbb{P}(E_U^C) + \mathbb{P}(E_\epsilon^C) \leq \delta_H + \delta_\epsilon$. If $E_G$ holds, the policy ensures that the true generator $\Gamma_B$ is continuously bounded by the required rate $R_k := \big(1 - B(\boldsymbol{x}_k) - \Delta_\rho\big)/T$ for all $t \in [t_k, t_k + T]$. This satisfies the underlying RA-CBF condition (e.g., Prop. III.2 or IV.1), so $\mathbb{P}(F \mid E_G) \leq \rho_d$. Using the law of total probability: $\mathbb{P}(F) = \mathbb{P}(F \mid E_G)\mathbb{P}(E_G) + \mathbb{P}(F \mid E_G^C)\mathbb{P}(E_G^C) \leq \rho_d(1 - \mathbb{P}(E_G^C)) + 1 \cdot \mathbb{P}(E_G^C)$. Substituting the bound for $\mathbb{P}(E_G^C)$ yields the result: $\mathbb{P}(F) \leq \rho_d + (1-\rho_d)(\delta_H + \delta_\epsilon) = \rho_d + \delta_H + \delta_\epsilon - \rho_d(\delta_H + \delta_\epsilon)$. □

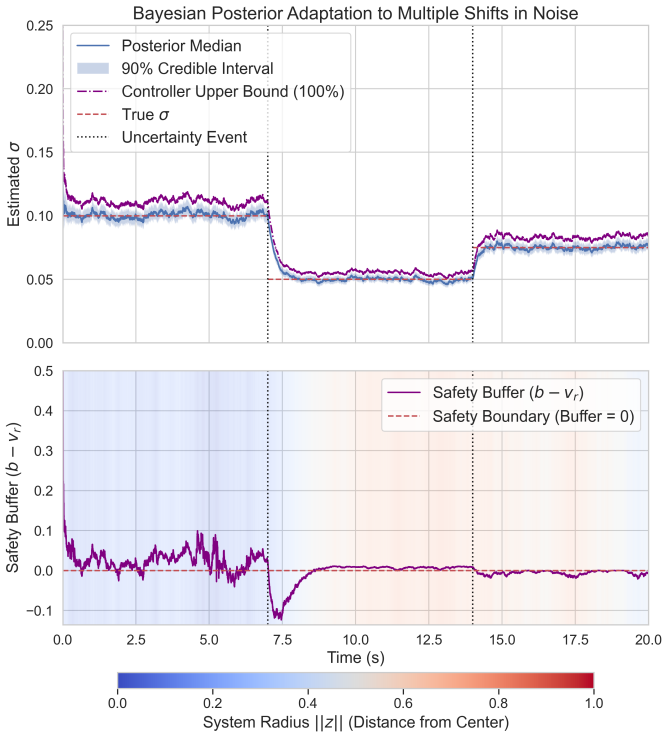


Fig. 2. Bayesian RA-CBF adaptation for the single-integrator with circular barrier $B(\boldsymbol{x}) = \|\boldsymbol{x}\|^2 / R_c^2$ ($R_c = 1$). **Top:** Posterior for the diffusion scale $\sigma$ (median and 90% credible band) under unannounced step changes in the true value (red dashed line). **Bottom:** Safety buffer ($\mathfrak{b} - v_r$), where $\mathfrak{b}$ is the maximum outward radial speed allowed by the Bayesian RA-CBF constraint and $v_r$ is the realized outward radial speed. A nonnegative buffer means the applied control remains inside the admissible risk-aware envelope. The background shows the normalized radius $\|\boldsymbol{x}_t\|/R_c$.

### C. QP-based Controller Synthesis

The constraint (18) is affine in $\boldsymbol{u}$. We formulate the controller at $t_k$ to minimally modify a nominal controller $\boldsymbol{u}_0(\boldsymbol{x}_k)$, introducing a slack $s$ for feasibility:

$$(\boldsymbol{u}_k^\star, s_k^\star) = \arg \min_{\boldsymbol{u} \in \mathcal{U},\ s \geq 0} \tfrac{1}{2}\|\boldsymbol{u} - \boldsymbol{u}_0(\boldsymbol{x}_k)\|^2 + \tfrac{1}{2} w s^2, \tag{22a}$$

$$\text{s.t.} \quad L_g B(\boldsymbol{x}_k)\, \boldsymbol{u} \leq \frac{1 - B(\boldsymbol{x}_k) - \Delta_\rho}{T} - \mu_\Gamma(\boldsymbol{x}_k, 0) - q_{\delta_H/N} \sigma_\Gamma(\boldsymbol{x}_k) - \epsilon(\Delta t, \delta_\epsilon/N) + s. \tag{22b}$$

The slack variable $s \geq 0$, penalized in the objective by a large weight $w \gg 0$, ensures the QP is always feasible, even if no control input $\boldsymbol{u}$ can satisfy the safety condition. However, the formal probabilistic safety guarantees are predicated on the constraint in (22b) being satisfied, which requires the optimal slack to be zero ($s_k^\star = 0$). The online implementation is summarized in Algorithm 1.

### D. Special Case: Bayesian Calibration of Noise Sensitivity

When posterior uncertainty is concentrated in the scalar noise proxy $\eta$, propagating the full generator posterior can be unnecessary. In that case, part of the total risk budget can be assigned to epistemic uncertainty in $\eta$, with the remainder reserved for aleatoric failure conditioned on a conservative

quantile of $p(\eta \mid \mathcal{D})$. We therefore design the controller using a high-probability quantile of the posterior $p(\eta \mid \mathcal{D})$ and adjust the internal risk target accordingly.

**Proposition V.6** (Risk Calibration)**.** Let $\rho_{\text{req}}$ be the required total risk. Let $\bar{\eta}$ be the $(1-\delta)$ quantile of the posterior of $\eta$. If a controller is designed using $\bar{\eta}$ with internal risk $\rho_d$, such that $\mathbb{P}(\text{failure} \mid \eta \le \bar{\eta}) \le \rho_d$, then setting

$$\rho_d = \max\left\{0,\ \frac{\rho_{\text{req}} - \delta}{1-\delta}\right\}$$

ensures $\mathbb{P}(\text{failure}) \le \rho_{\text{req}}$.

*Proof.* Let $F$ be the failure event. By the law of total probability:

$$\begin{aligned}\mathbb{P}(F) &= \mathbb{P}(F \mid \eta \le \bar{\eta})\mathbb{P}(\eta \le \bar{\eta}) + \mathbb{P}(F \mid \eta > \bar{\eta})\mathbb{P}(\eta > \bar{\eta}) \\ &\le \rho_d(1-\delta) + 1 \cdot \delta.\end{aligned}$$

Set $\mathbb{P}(F) \le \rho_{\text{req}}$ and solve for $\rho_d$. □

**Algorithm 1** Online Bayesian RA-CBF (RHC)

Input: $\rho_d$, $T = N\Delta t$, $\delta_H, \delta_\epsilon, \Delta t$.
Initialize GP priors for $f, D$; collect seed data $\mathcal{D}_0$.
Offline: compute conservative $\eta$ and $\epsilon(\Delta t, \delta_\epsilon/N)$.
**for** $k = 0, 1, 2, \ldots$ **do**
    Update posteriors $\mu_f, \Sigma_f,\ \mu_D, \Sigma_D$ at $\boldsymbol{x}_k$.
    Compute $\mu_\Gamma(\boldsymbol{x}_k, 0), \sigma_\Gamma(\boldsymbol{x}_k)$ via (16).
    Form QP constraint (22b) using $q_{\delta_H/N}, \epsilon$.
    Solve QP (22) for $\boldsymbol{u}_k^\star$; apply to system.
    Observe $\boldsymbol{x}_{k+1}$, update dataset $\mathcal{D}_{k+1}$.
**end for**

## VI. Conclusion

This work introduces a Bayesian Risk-Aware Control Barrier Function (RA-CBF) framework for discrete-time stochastic systems with learned dynamics, providing a robust and practical approach to safety in uncertain environments. We establish two distinct discrete-time safety guarantees: a robust, time-uniform bound via martingale concentration and a tighter pathwise bound that recovers the non-conservative performance of continuous-time designs. To manage model uncertainty, our Bayesian approach propagates posterior uncertainty over the system's drift and diffusion into a high-probability confidence buffer and an Upper Confidence Bound (UCB) for the generator. This formulation, combined with an inter-sample robustness margin, yields a high-confidence probabilistic safety guarantee for receding horizon control, implemented as an efficient, convex Quadratic Program (QP) that computes minimally invasive actions respecting both aleatoric risk and epistemic uncertainty.

## References


[1] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs for safety critical systems," *IEEE Trans. on Automatic Control*, vol. 62, no. 8, pp. 3861–3876, 2017.

[2] Y. Chen, A. Singletary, and A. D. Ames, "Guaranteed obstacle avoidance for multi-robot operations with limited actuation: A control barrier function approach," *IEEE Control Systems Letters*, vol. 5, no. 1, pp. 127–132, 2021.

[3] W. Shaw Cortez, D. Oetomo, C. Manzie, and P. Choong, "Control barrier functions for mechanical systems: Theory and application to robotic grasping," *IEEE Transactions on Control Systems Technology*, vol. 29, no. 2, pp. 530–545, 2021.

[4] K. Garg and D. Panagou, "Robust control barrier and control lyapunov functions with fixed-time convergence guarantees," in *2021 American Control Conference (ACC)*, 2021, pp. 2292–2297.

[5] P. Jagtap, G. J. Pappas, and M. Zamani, "Control barrier functions for unknown nonlinear systems using gaussian processes," in *59th IEEE Conference on Decision and Control*, 2020, pp. 3699–3704.

[6] M. Aali and J. Liu, "Learning high-order control barrier functions for safety-critical control with gaussian processes," in *2024 American Control Conference (ACC)*. IEEE, 2024, pp. 1–6.

[7] R. Gutierrez and J. B. Hoagg, "Control barrier functions with real-time gaussian process modeling," *arXiv preprint arXiv:2505.06765*, 2025.

[8] K. Long, Y. Yi, Z. Dai, S. Herbert, J. Cortés, and N. Atanasov, "Sensor-based distributionally robust control for safe robot navigation in dynamic environments," *The International Journal of Robotics Research*, p. 02783649251352000, 2025.

[9] A. Chriat and C. Sun, "Wasserstein distributionally robust control barrier function using conditional value-at-risk with differentiable convex programming," in *AIAA SCITECH 2024 Forum*, 2024, p. 0725.

[10] A. Hakobyan and I. Yang, "Wasserstein distributionally robust motion control for collision avoidance using conditional value-at-risk," *IEEE Transactions on Robotics*, vol. 38, no. 2, pp. 939–957, 2021.

[11] A. Baheri, "Distributionally robust lyapunov–barrier networks for safe and stable control under uncertainty," *Results in Control and Optimization*, vol. 19, p. 100556, 2025.

[12] C. Santoyo, M. Dutreix, and S. Coogan, "A barrier function approach to finite-time stochastic system verification and control," *Automatica*, vol. 125, p. 109439, 2021.

[13] S. Yaghoubi, K. Majd, G. Fainekos, T. Yamaguchi, D. Prokhorov, and B. Hoxha, "Risk-bounded control using stochastic barrier functions," *IEEE Control Systems Letters*, vol. 5, no. 5, pp. 1831–1836, 2021.

[14] M. Black, G. Fainekos, B. Hoxha, D. Prokhorov, and D. Panagou, "Safety under uncertainty: Tight bounds with risk-aware control barrier functions," *arXiv preprint arXiv:2304.01040*, 2023.

[15] J. Breeden, K. Garg, and D. Panagou, "Control barrier functions in sampled-data systems," *IEEE Control Systems Letters*, vol. 6, pp. 367–372, 2022.

[16] M. Vahs, C. Pek, and J. Tumova, "Belief control barrier functions for risk-aware control," *IEEE Robotics and Automation Letters*, vol. 8, no. 12, pp. 8565–8572, 2023.

[17] G. Bahati, P. Ong, and A. D. Ames, "Sample-and-hold safety with control barrier functions," in *2024 American Control Conference (ACC)*, 2024, pp. 5169–5176.

[18] J. Li, Q. Liu, W. Jin, J. Qin, and S. Hirche, "Robust safe learning and control in an unknown environment: An uncertainty-separated control barrier function approach," *IEEE Robotics and Automation Letters*, vol. 8, no. 10, pp. 6539–6546, 2023.

[19] R. K. Cosner, P. Culbertson, and A. D. Ames, "Bounding stochastic safety: Leveraging freedman's inequality with discrete-time control barrier functions," *IEEE Control Systems Letters*, vol. 8, pp. 1937–1942, 2024.

[20] F. Shmarov, S. Soudjani, N. Paoletti, E. Bartocci, S. Lin, S. A. Smolka, and P. Zuliani, "Automated synthesis of safe digital controllers for sampled-data stochastic nonlinear systems," *IEEE Access*, vol. 8, pp. 180 825–180 843, 2020.

[21] D. Revuz and M. Yor, *Continuous martingales and Brownian motion*. Springer Science & Business Media, 2013, vol. 293.

[22] P. Mörters and Y. Peres, *Brownian motion*. Cambridge University Press, 2010, vol. 30.

[23] J.-F. Le Gall, *Brownian motion, martingales, and stochastic calculus*. Springer, 2016.

[24] S. R. Howard, A. Ramdas, J. McAuliffe, and J. Sekhon, "Time-uniform Chernoff bounds via nonnegative supermartingales," *Probability Surveys*, vol. 17, no. none, pp. 257 – 317, 2020. [Online]. Available: https://doi.org/10.1214/18-PS321

[25] C. K. Williams and C. E. Rasmussen, *Gaussian processes for machine learning*. MIT press Cambridge, MA, 2006, vol. 2, no. 3.

[26] R. Vershynin, *High-dimensional probability: An introduction with applications in data science*. Cambridge university press, 2018, vol. 47.