

# Control Design for Risk-Based Signal Temporal Logic Specifications

Sleiman Safaoui, Lars Lindemann, Dimos V. Dimarogonas, Iman Shames, Tyler H. Summers

**Abstract**—We present a general framework for risk semantics on Signal Temporal Logic (STL) specifications for stochastic dynamical systems using axiomatic risk theory. We show that under our recursive risk semantics, risk constraints on STL formulas can be expressed in terms of risk constraints on atomic predicates. We then show how this allows a (*stochastic*) STL risk constraint to be transformed into a risk-tightened *deterministic* STL constraint on a related deterministic nominal system, enabling the application of existing STL methods. For affine predicate functions and a (coherent) Distributionally Robust Value at Risk measure, we show how risk constraints on atomic predicates can be reformulated as tightened deterministic affine constraints. We demonstrate the framework using a Model Predictive Control (MPC) design with an STL risk constraint.

**Index Terms**—Stochastic systems, optimization, constrained control.

## I. INTRODUCTION

TEMPORAL logics allow to reason about temporal properties of systems and have traditionally been used in formal verification and model checking [1]. More recently, temporal logics have also been used to impose highly expressive mission specifications on complex autonomous systems. For systems under linear temporal logic (LTL) and metric interval temporal logic (MITL) specifications, motion planning and control synthesis algorithms have been proposed in [2]–[4].

*Signal Temporal Logic (STL)* is a temporal logic interpreted over dense-time real-valued (*deterministic*) signals [5] similar to MITL; it allows to additionally impose quantitative spatial properties by means of predicates that go beyond the abstract use of propositions in LTL and MITL. STL is hence more expressive and has been the focus of motion planning and control synthesis in areas such as robotics. In addition to the Boolean satisfaction relation given for LTL and MITL specifications, one can associate quantitative semantics with an STL specification that allow to reason about how robustly (severely) a specification is satisfied (violated). These quantitative semantics come in the form of the robustness degree

and the robust semantics [6] as well as space robustness, time robustness, and other variants [7]–[9].

Control of systems under STL specifications is inherently different compared to systems under LTL and MITL specifications, which is mainly based on abstractions and automata theory. For deterministic discrete-time systems, the authors in [10] transform the STL specification into mixed-integer linear constraints and use Model Predictive Control (MPC) to deal with these constraints. Similarly, MPC has been employed by maximizing certain forms of the quantitative semantics associated with the STL specification at hand [8], [9], [11]. These methods seem computationally more tractable than the approach presented in [10] due to the use of smooth quantitative semantics. For deterministic continuous-time systems, [12] uses timed-automata theory to decompose the STL specification into STL subspecifications that can be implemented by low-level feedback control laws, such as those in [13]–[15]. Learning-based methods for partially unknown systems have been presented in [16]–[18].

A large fraction of the aforementioned research focuses on finite abstractions and/or *deterministic* systems. However, methodological advances are required to account for inherent uncertainties and high-dimensional continuous spaces in autonomous systems, especially due to *stochastic uncertainties* arising from the use of noisy data and learning components. Robust extensions of [10] have been presented in [19], [20]. Probabilistic notions of STL for stochastic systems have been presented in [21], [22]. However, these formulations are either deterministic and based on a worst-case approach or utilize chance constraints, an incoherent measure of risk that can lead to undesirable decisions. Effective risk management in complex autonomous systems demands a more sophisticated approach to quantifying risks of specification violations. This motivates an axiomatic approach to risk, which has been advocated for in finance [23] and more recently in robotics [24]. Control under coherent risk measures has been considered in [25], [26]. The extension to more complex and generic specifications, such as those captured by temporal logics, has however not been addressed.

**Contributions.** 1) We present a general framework for defining risk semantics of STL specifications for *stochastic dynamical systems* using axiomatic risk theory. In particular, we compose risk metrics with predicate functions, which become stochastic in the considered setup. 2) We then recursively define risk semantics for Boolean and temporal STL operators. For a given STL specification, we show that these risk semantics can be expressed as risk constraints on predicate functions over certain time intervals. We then show

This material is based on work supported by the United States Air Force Office of Scientific Research under award number FA2386-19-1-4073. This work was also supported in part by the Swedish Research Council (VR) and Defence Science and Technology Group, Australian Government, through agreement MyIP: ID9156.

S. Safaoui is with the department of Electrical Engineering and T. H. Summers is with the department of Mechanical Engineering at University of Texas at Dallas, Richardson, TX, USA, E-mail: {sleiman.safaoui, tyler.summers}@utdallas.edu. L. Lindemann and D. V. Dimarogonas are with the Division of Decision and Control Systems, KTH Royal Institute of Technology, 100 44 Stockholm, Sweden, E-mail: {llindem, dimos}@kth.se. I. Shames is with the University of Melbourne, Parkville, VIC 3010, Australia, E-mail: iman.shames@unimelb.edu.au.

how this allows such a risk constraint to be transformed into a risk-tightened deterministic STL constraint on a related deterministic nominal system. 3) For affine predicates and a coherent Distributionally Robust Value at Risk measure, we show how risk constraints on predicate functions can be explicitly reformulated as tightened deterministic affine constraints. 4) To demonstrate the framework, we use an MPC formulation to solve a risk-based STL control design problem, which we illustrate with numerical experiments.

## II. RISK MEASURES AND AXIOMATIC RISK THEORY

A risk measure is a function that assigns a real number to a random variable, quantifying its size, typically related to one of its tails. More formally, let  $(\Omega, \mathcal{F}, \mathbb{P})$  be a probability space, where  $\Omega$  is the sample space,  $\mathcal{F}$  is a  $\sigma$ -algebra of subsets of  $\Omega$ , and  $\mathbb{P}$  is a probability measure on  $\mathcal{F}$ . Let  $\mathbf{X}$  denote a set of real-valued random variables on  $\Omega$  (i.e.,  $X \in \mathbf{X}$  is a Borel measurable function  $X : \Omega \rightarrow \mathbb{R}$ ). A risk measure is a function  $\rho : \mathbf{X} \rightarrow \mathbb{R} \cup \{\pm\infty\}$ . In finance, elements of  $\mathbf{X}$  represent the value of a financial position, and a risk measure quantifies the probability and severity of a financial loss. Here, elements of  $\mathbf{X}$  will represent states of a stochastic system, and a *risk measure quantifies the probability and severity of violating a signal temporal logic specification*.

**Risk Axioms.** Effective quantitative risk management in emerging complex autonomous systems is a major challenge, which motivates an axiomatic approach to risk measures. Four important axioms for a risk measure  $\rho$  [24] are:

- 1) **Monotonicity:** If  $X_1 \leq X_2$  then  $\rho(X_1) \leq \rho(X_2)$ ;
- 2) **Translation Invariance:**  $\rho(X + c) = \rho(X) + c$ ;
- 3) **Positive Homogeneity:**  $\rho(\beta X) = \beta \rho(X)$ ;
- 4) **Subadditivity:**  $\rho(X_1 + X_2) \leq \rho(X_1) + \rho(X_2)$ ;

$\forall X, X_1, X_2 \in \mathbf{X}$ , and  $c, \beta \in \mathbb{R}, \beta \geq 0$ . A risk measure is called *coherent* if it satisfies all four of these axioms [23]. It has been argued that these axioms constitute natural desirable properties for risk measures in complex systems [23], [24]. *Spectral or distortion risk measures* also satisfy

- 5) **Law Invariance:** If  $X_1, X_2 \in \mathbf{X}$  are identically distributed, then  $\rho(X_1) = \rho(X_2)$ .
- 6) **Comonotone Additivity:** If  $X_1, X_2 \in \mathbf{X}$  are comonotone (namely,  $(X_1(\omega_2) - X_1(\omega_1))(X_2(\omega_2) - X_2(\omega_1)) \geq 0 \forall \omega_1, \omega_2 \in \Omega$ ), then  $\rho(X_1 + X_2) = \rho(X_1) + \rho(X_2)$

and can be viewed as refinements of coherent risk measures.

**Examples of Risk Measures.** Several commonly used and studied risk measures operating on  $X \in \mathbf{X}$  include:

- 1) **Expectation:**  $\rho_E(X) = \mathbb{E}(X)$
- 2) **Worst-Case:**  $\rho_W(X) = \text{ess sup}(X)$
- 3) **Mean-Variance:**  $MV_\lambda(X) = \mathbb{E}(X) + \lambda \text{Var}(X)$ ,  $\lambda > 0$
- 4) **Value at Risk (VaR) at level  $\delta \in (0, 1]$ :**  $\text{VaR}_\delta(X) = \inf_{t \in \mathbb{R}} \{t \mid \mathbb{P}(X \leq t) \geq 1 - \delta\}$
- 5) **Conditional Value at Risk (CVaR) at level  $\delta \in (0, 1]$ :**  $\text{CVaR}_\delta(X) = \inf_{t \in \mathbb{R}} \{t + \frac{1}{\delta} \mathbb{E}[X - t]_+\}$ ,  $[\cdot]_+ := \max(\cdot, 0)$
- 6) **Entropic Value at Risk (EVAR) at level  $\delta \in (0, 1]$ :**  $\text{EVAR}_\delta(X) = \inf_{z > 0} \{z^{-1} \ln(M_X(z)/\alpha)\}$ , where  $M_X(z) = \mathbb{E}(e^{zX})$  is the moment generating function of  $X$ .

Unfortunately, many widely used risk measures in robotics and engineering are not coherent, and can lead to serious miscalculations of risk, e.g., mean-variance and mean-standard-deviation fail to be monotone, and VaR lacks subadditivity. The widely used chance constraint in optimization models is closely related to VaR and has been used for notions of STL robustness for stochastic systems. We advocate for axiomatic risk theory with coherent risk as a more systematic and sophisticated approach to risk management for STL specifications in emerging safety-critical autonomous systems.

**Distributional Robustness.** Evaluating any of the above risk measures requires knowledge of the probability distribution of the associated random variable. In practice however, we are never given the probability distribution, only noisy data. Instead, we must estimate properties of the distribution from the noisy data, or make assumptions about the distribution. In the emerging area of distributionally robust optimization [27], this uncertainty in our knowledge of the probability distribution itself is explicitly accounted for. Rather than assuming a single probability distribution, we instead work with *ambiguity sets* of distributions. Common examples of ambiguity sets include those with certain moment constraints, e.g., the set of distributions with given mean and variance  $\mathcal{P} = \{\mathbb{P} \mid \mathbb{E}_{\mathbb{P}}[X] = \mu, \mathbb{E}_{\mathbb{P}}[(X - \mu)^2] = \Sigma\}$ , or distance/divergence based sets, such as  $\mathcal{P} = \{\mathbb{P} \mid d(\mathbb{P}, \hat{\mathbb{P}}) \leq \varepsilon\}$ , where  $d$  is a probability distance function,  $\hat{\mathbb{P}}$  is a nominal distribution, and  $\varepsilon > 0$  is a radius. A risk measure and an ambiguity set can be combined to obtain distributionally robust (DR) risk measures:  $\rho_{\text{DR}} = \sup_{\mathbb{P} \in \mathcal{P}} \rho(X)$ . For example, DR-VaR,  $\sup_{\mathbb{P} \in \mathcal{P}} \text{VaR}_\delta(X)$ , with various moment-based ambiguity sets is a coherent risk measure [28].

## III. RISK-BASED SIGNAL TEMPORAL LOGIC

Traditionally, STL constraints are specified for *deterministic* dynamical systems. STL formulas are based on predicates  $\pi \in \{\top, \perp\}$  where  $\top$  and  $\perp$  denote true and false, respectively, and are obtained from evaluating a predicate function  $\alpha : \mathbb{R}^n \rightarrow \mathbb{R}$  such that  $\pi = \{\alpha(x) \geq 0\}$ , i.e.,  $\pi = \top$  if and only if  $\alpha(x) \geq 0$  where  $x \in \mathbb{R}$ . A typical STL formula  $\varphi$  is composed from logical and bounded-time temporal operators and can always be rewritten in *negation normal form* [6] (also referred to as positive normal form in [19]), where the negation operator appears only at the atomic predicate level. We assume that STL formulas are given in negation normal form, which are defined recursively through the grammar:

$$\varphi := \top \mid \pi \mid \neg \pi \mid \varphi' \wedge \varphi'' \mid \varphi' \vee \varphi'' \mid \varphi' \mathcal{U}_{[a,b]} \varphi'' \mid \varphi' \mathcal{R}_{[a,b]} \varphi'' \quad (1)$$

where  $\neg$  is the negation operator,  $\wedge$  and  $\vee$  are conjunction and disjunction,  $\mathcal{U}_{[a,b]}$  and  $\mathcal{R}_{[a,b]}$  are the until and release operators with  $a \leq b$  and  $a, b \in \mathbb{R}_{\geq 0} = [0, +\infty)$ . Note that  $\vee$  and  $\mathcal{R}$  are duals of  $\wedge$  and  $\mathcal{U}$  respectively, since  $\varphi' \wedge \varphi'' = \neg(\neg \varphi' \vee \neg \varphi'')$  and  $\varphi' \mathcal{U}_{[a,b]} \varphi'' = \neg(\neg \varphi' \mathcal{R}_{[a,b]} \neg \varphi'')$ . The negation normal form replaces the negation of a formula by including all operators and their duals in the grammar. We can also define additional operators such as eventually ( $\diamond_{[a,b]} \varphi := \top \mathcal{U}_{[a,b]} \varphi$ ) and always ( $\square_{[a,b]} \varphi := \perp \mathcal{R}_{[a,b]} \varphi$ ).

Denote by  $\xi_N$  a *deterministic* discrete-time finite-horizon system run of length  $N + 1$ . If the signal starts at time

$t$ , then  $(\xi_N, t) = x_t x_{t+1} \dots x_{t+N}$  where  $x_i$  is the *deterministic* state at time  $i$ .  $(\xi_N, t) \models \varphi$  denotes  $\xi_N$  satisfying  $\varphi$  (starting) at time  $t$ . By convention,  $\xi_N$  satisfies  $\varphi$  (denoted  $\xi_N \models \varphi$ ) if  $(\xi_N, 0) \models \varphi$ . The traditional STL semantics are defined as:  $(\xi_N, t) \models \top$  trivially holds,  $(\xi_N, t) \models \pi$  iff  $\alpha(x_t) \geq 0$ ,  $(\xi_N, t) \models \neg\pi$  iff  $\neg((\xi_N, t) \models \pi)$ ,  $(\xi_N, t) \models \varphi \wedge \varphi'$  iff  $(\xi_N, t) \models \varphi \wedge (\xi_N, t) \models \varphi'$ ,  $(\xi_N, t) \models \varphi \vee \varphi'$  iff  $(\xi_N, t) \models \varphi \vee (\xi_N, t) \models \varphi'$ ,  $(\xi_N, t) \models \varphi \mathcal{U}_{[a,b]} \varphi'$  iff  $\exists t' \in [t+a, t+b]$  s.t.  $(\xi_N, t') \models \varphi' \wedge \forall t'' \in [t, t']$ ,  $(\xi_N, t'') \models \varphi$ , and  $(\xi_N, t) \models \varphi \mathcal{R}_{[a,b]} \varphi'$  iff  $\forall t' \in [t+a, t+b]$ ,  $(\xi_N, t') \models \varphi' \vee \exists t'' \in [t, t']$ ,  $(\xi_N, t'') \models \varphi$ . An STL formula horizon can be upper bounded by the maximum over the sums of the nested upper bounds of the temporal operators. The bound is denoted  $\text{len}(\varphi)$  and defined recursively as:  $\text{len}(\top) = \text{len}(\pi) = \text{len}(\neg\pi) = 0$ ,  $\text{len}(\varphi \wedge \varphi') = \text{len}(\varphi \vee \varphi') = \max(\text{len}(\varphi), \text{len}(\varphi'))$ , and  $\text{len}(\varphi \mathcal{U}_{[a,b]} \varphi') = \text{len}(\varphi \mathcal{R}_{[a,b]} \varphi') = b + \max(\text{len}(\varphi), \text{len}(\varphi'))$ . As opposed to the Boolean nature of the above semantics, STL allows to define a quantitative robustness measure, which is a function  $\rho^\varphi(\xi_N, t)$  [7] that assigns the formula  $\varphi$  a real value such that  $(\xi_N, t) \models \varphi$  if  $\rho^\varphi(\xi_N, t) > 0$ , indicating how robustly the formula is satisfied. Note that  $\rho^\varphi(\xi_N, t)$  is deterministic.

*Stochastic* systems require an alternative quantitative semantics. We denote a discrete-time finite-horizon *stochastic process* starting at time  $t$  as  $(\Xi_N, t) = X_t X_{t+1} \dots X_{t+N}$  where  $X_i$  is an  $\mathbb{R}^n$ -valued random variable from a set  $\mathbf{X}$  of random variables that represents the *stochastic* system state at time  $i$ . Since the system state is stochastic, the question of  $\Xi_N$  satisfying an STL formula is ill-posed, and the aforementioned Boolean semantics do not apply (notice that  $\alpha(X)$  is a random variable). Instead, we quantify the risk of *violating* a specification by introducing a risk measure into STL formulas. (This contrasts with STL robustness measures, which quantify the *satisfaction* of a formula; it is more natural to consider the risk of violating a formula.) While there are several ways to define the risk of *violating* a formula  $\mathcal{R}$ , we choose to define the risk of violating an atomic predicate using the risk measures  $\rho$  discussed in Section II and build up the STL risk-of-violation semantics recursively.

**Definition 1.** *STL Risk Semantics.* For the stochastic system run  $\Xi_N$ , risk measure  $\rho$ , atomic predicate  $\pi$ , and STL formulas  $\varphi$  and  $\varphi'$ , the STL risk is defined recursively as

- $\mathcal{R}((\Xi_N, t) \not\models \top) := -\infty$ ,
- $\mathcal{R}((\Xi_N, t) \not\models \perp) := +\infty$ ,
- $\mathcal{R}((\Xi_N, t) \not\models \pi) := \rho(-\alpha(X_t))$ ,
- $\mathcal{R}((\Xi_N, t) \not\models \neg\pi) := \rho(\alpha(X_t))$ ,
- $\mathcal{R}((\Xi_N, t) \not\models \varphi \wedge \varphi') := \max(\mathcal{R}((\Xi_N, t) \not\models \varphi), \mathcal{R}((\Xi_N, t) \not\models \varphi'))$ ,
- $\mathcal{R}((\Xi_N, t) \not\models \varphi \vee \varphi') := \min(\mathcal{R}((\Xi_N, t) \not\models \varphi), \mathcal{R}((\Xi_N, t) \not\models \varphi'))$ ,
- $\mathcal{R}((\Xi_N, t) \not\models \varphi \mathcal{U}_{[a,b]} \varphi') := \min_{i \in [a,b]} (\max(\mathcal{R}((\Xi_N, t+i) \not\models \varphi'), \max_{j \in [0,i]} \mathcal{R}((\Xi_N, t+j) \not\models \varphi)))$ ,
- $\mathcal{R}((\Xi_N, t) \not\models \varphi \mathcal{R}_{[a,b]} \varphi') := \max_{i \in [a,b]} (\min(\mathcal{R}((\Xi_N, t+i) \not\models \varphi'), \min_{j \in [0,i]} \mathcal{R}((\Xi_N, t+j) \not\models \varphi)))$ .

Often it is desirable to bound the STL risk of violation of a formula  $\varphi$  by some prescribed value  $\delta$ , i.e.,  $\mathcal{R}((\Xi_N, t) \not\models \varphi) \leq \delta$ . The risk metrics of Section II can be used to quantify

the aforementioned risk. It must be noted that these metrics are generally asymmetric, i.e.  $\rho(-\alpha(X_t)) \neq -\rho(\alpha(X_t))$ , and appear only at the atomic level (which further justifies using the grammar (1)). This will be explored in Section V.

#### IV. CONTROL DESIGN FOR RISK-CONSTRAINED STL

We consider the discrete-time stochastic linear system:

$$X_{t+1} = A_t X_t + B_t u_t + W_t, \quad X_0 = x_0, \quad (2)$$

where  $X_t \in \mathbb{R}^n$ ,  $u_t \in \mathbb{R}^m$ , and  $W_t \in \mathbb{R}^n$  are respectively the state, control input, and stochastic disturbance of the system at time  $t$ ,  $A_t \in \mathbb{R}^{n \times n}$  and  $B_t \in \mathbb{R}^{n \times m}$  are respectively the system dynamics and input matrices, and  $x_0$  is the known initial system state. The disturbances  $\{W_t\}$  are assumed independent according to unknown distributions  $\mathbb{P}_{W_t}$ , which belong to ambiguity sets  $\mathcal{P}_t^W$ . Here we consider moment-based ambiguity sets with given mean and variance:  $\mathcal{P}_t^W = \{\mathbb{P}_{W_t} \mid \mathbb{E}[W_t] = \bar{W}_t, \mathbb{E}[(W_t - \bar{W}_t)(W_t - \bar{W}_t)^T] = \Sigma_{W_t}\}$ .

From (2) we obtain

$$X_\tau = \Phi(\tau, t) X_t + \sum_{k=1}^{\tau-1} \Phi(\tau, k+1) (B_k u_k + W_k), \quad (3)$$

where  $\Phi(\cdot, \cdot)$  is the state transition matrix defined as

$$\Phi(\tau, t) = \begin{cases} A_{\tau-1} A_{\tau-2} \dots A_t & \tau > t \geq 0 \\ I_n & \tau = t \geq 0. \end{cases} \quad (4)$$

We consider a finite-horizon control design problem where the goal is to determine a state feedback control policy for (2) that satisfies a risk constraint associated with an STL formula  $\varphi$ . Specifically, we consider

$$J^*(x_0) = \min_{\mu \in \mathcal{M}} \mathbb{E}[J(X_{0:N}, u_{0:N-1})] \quad \text{subject to} \quad (5)$$

$$X_t = \Phi(t, 0) x_0 + \sum_{k=0}^{t-1} \Phi(t, k+1) (B_k u_k + W_k),$$

$$\mathcal{R}((\Xi_N, 0) \not\models \varphi) \leq \delta, \quad u_{0:N-1} \in U,$$

where the decision variable is a state feedback policy  $\mu = [\mu_0, \mu_1, \dots, \mu_{N-1}]$  with  $u_t = \mu_t(x_t)$  from a set  $\mathcal{M}$  of measurable policies,  $u_{t:N-1} = [u_t^T, u_{t+1}^T, \dots, u_{N-1}^T]^T$ ,  $X_{t:N} = [X_t^T, X_{t+1}^T, \dots, X_{N-1}^T]^T$  and  $W_{t:N-1} = [W_t^T, W_{t+1}^T, \dots, W_{N-1}^T]^T$ ,  $N$  is the time horizon,  $J$  is a stage cost function with expectation taken with respect to the disturbance sequence  $\{W_t\}$ ,  $\delta \in \mathbb{R}$  is a user-defined risk bound, and  $U \subset \mathbb{R}^{Nm}$  is an input constraint set. The challenge lies in the uncertainty in the system model and its appearance in the STL risk constraint. We approach this using Model Predictive Control (MPC) and a reformulation of the STL risk constraint into a tightened deterministic STL constraint so that existing control approaches for deterministic STL constraints can be used such as summarized in Section I.

Using MPC to (approximately) solve (5) reduces the problem to a sequence of open-loop optimization problems. For an STL problem with formula  $\varphi$ , a natural choice for the prediction horizon is  $N \geq \text{len}(\varphi)$  with a system run of  $N_s > N$ . The MPC optimization problem at time  $t \leq N_s - N$  given the observed state,  $x_t$ , is:

$$J^*(x_t) = \min_{u_{t:t+N-1}} \mathbb{E}[J(X_{t:t+N}, u_{t:t+N-1})] \quad \text{subject to} \quad (6)$$

$$X_\tau = \Phi(\tau, t) x_t + \sum_{k=t}^{\tau-1} \Phi(\tau, k+1) (B_k u_k + W_k),$$

$$\mathcal{R}((\Xi_N, 0) \not\# \varphi) \leq \delta, \quad u_{t:t+N-1} \in U,$$

where the decision variable is the open-loop control sequence  $u_{t:t+N-1}$ , expectation is with respect to  $W_{t:t+N-1}$ . Solving the optimization problem yields the future optimal control sequence  $u_{t:t+N-1}^* = [u_t^*, \dots, u_{t+N-1}^*]$ . Only the first component  $u_t^*$  of this plan is implemented, and the problem is solved again after the next state realization is observed. Thus, the MPC policy is  $\mu_{\text{MPC}}(x_t) = u_t^*$ . We assume that  $\mathbb{E}[J(X_{t:t+N}, u_{t:t+N-1})]$  in (6) can be evaluated analytically, which is the case for quadratic  $J$ , allowing us to focus solely on challenges in accounting for the STL risk constraint.

## V. REFORMULATION OF STL RISK CONSTRAINTS

In this section, we demonstrate how the optimization problem (6) can be reformulated into an optimization problem with *deterministic*, tightened STL constraints on the nominal system dynamics. We also show how to explicitly write the tightened constraints on atomic predicates for one specific coherent risk measure, Distributionally Robust Value at Risk.

### A. From STL formula violation risk to atomic predicate violation risk

The semantics in Definition 1 are useful for two main reasons: 1) the risk metrics described in Section II appear only at the atomic predicate level and the risk of failing to satisfy an STL formula is defined recursively from there, 2) through the recursive STL risk semantics, all operators (except negation) are defined in terms of min and max operators over time intervals. These reasons allow transforming a risk-based STL constraint into similar risk constraints on atomic predicates. This is formalized in Theorem 1.

**Theorem 1.** Consider a risk measure  $\rho$  from Section II, STL risk semantics  $\mathcal{R}$  from Definition 1, STL formula  $\varphi$  given by (1), and risk bound  $\delta \in \mathbb{R}$ . The STL risk constraint  $\mathcal{R}((\Xi_N, t) \not\# \varphi) \leq \delta$  can be transformed into conjunctions and disjunctions over time intervals of risk constraints on atomic predicates of the form  $\rho(\pm\alpha(X_t)) \leq \delta$  where  $\alpha$  is a predicate function associated with the atomic predicate  $\pi$ .

*Proof.* The proof is a consequence of the grammar (1) and the structure of the risk semantics which apply the risk measure only on the atomic predicate level. We consider eight cases (other cases follow recursively):

- 1)  $\mathcal{R}((\Xi_N, t) \not\# \top) \leq \delta \Leftrightarrow -\infty \leq \delta$  (always true)
- 2)  $\mathcal{R}((\Xi_N, t) \not\# \perp) \leq \delta \Leftrightarrow +\infty \leq \delta$  (always false)
- 3)  $\mathcal{R}((\Xi_N, t) \not\# \pi) \leq \delta \Leftrightarrow \rho(-\alpha(X_t)) \leq \delta$
- 4)  $\mathcal{R}((\Xi_N, t) \not\# \neg\pi) \leq \delta \Leftrightarrow \rho(\alpha(X_t)) \leq \delta$
- 5)  $\mathcal{R}((\Xi_N, t) \not\# \varphi \wedge \varphi') \leq \delta$   
 $\Leftrightarrow \max(\mathcal{R}((\Xi_N, t) \not\# \varphi), \mathcal{R}((\Xi_N, t) \not\# \varphi')) \leq \delta$   
 $\Leftrightarrow \mathcal{R}((\Xi_N, t) \not\# \varphi) \leq \delta \wedge \mathcal{R}((\Xi_N, t) \not\# \varphi') \leq \delta$
- 6)  $\mathcal{R}((\Xi_N, t) \not\# \varphi \vee \varphi') \leq \delta$   
 $\Leftrightarrow \min(\mathcal{R}((\Xi_N, t) \not\# \varphi), \mathcal{R}((\Xi_N, t) \not\# \varphi')) \leq \delta$   
 $\Leftrightarrow \mathcal{R}((\Xi_N, t) \not\# \varphi) \leq \delta \vee \mathcal{R}((\Xi_N, t) \not\# \varphi') \leq \delta$
- 7)  $\mathcal{R}((\Xi_N, t) \not\# \varphi \mathcal{A}_{[a,b]}) \leq \delta$   
 $\Leftrightarrow \min_{i \in [a,b]} (\max(\mathcal{R}((\Xi_N, t+i) \not\# \varphi'),$   
 $\max_{j \in [0,i]} \mathcal{R}((\Xi_N, t+j) \not\# \varphi))) \leq \delta$

$$\begin{aligned} &\Leftrightarrow \exists t' \in [t+a, t+b] \mathcal{R}((\Xi_N, t') \not\# \varphi') \leq \delta \\ &\wedge \forall t'' \in [t, t'] \mathcal{R}((\Xi_N, t'') \not\# \varphi) \leq \delta \\ 8) &\mathcal{R}((\Xi_N, t) \not\# \varphi \mathcal{B}_{[a,b]}) \leq \delta \\ &\Leftrightarrow \max_{i \in [a,b]} (\min(\mathcal{R}((\Xi_N, t+i) \not\# \varphi'), \\ &\min_{j \in [0,i]} \mathcal{R}((\Xi_N, t+j) \not\# \varphi))) \leq \delta \\ &\Leftrightarrow \forall t' \in [t+a, t+b] \mathcal{R}((\Xi_N, t') \not\# \varphi') \leq \delta \\ &\vee \exists t'' \in [t, t'] \mathcal{R}((\Xi_N, t'') \not\# \varphi) \leq \delta \end{aligned}$$

The  $\forall$  and  $\exists$  operators can be expressed as chains of  $\wedge$  and  $\vee$  operators respectively across different time steps resulting in time intervals over which conjunctions and disjunctions of deterministic constraints (lines 3 and 4) are formed. Lines 1 and 2 are trivial and can be safely discarded.  $\square$

Using Theorem 1, it is clear that the risk level  $\delta$  is carried through the formula and only appears at the atomic level. All operators, other than negations, are also preserved and they reappear to connect the risk of violating atomic predicates; e.g. the risk of failing to satisfy the conjunction of two predicates is the conjunction of the risk of failing to satisfy each predicate. Furthermore, the risk measure of an atomic predicate integrates the stochasticity in the state, leaving us with *deterministic* inequalities of the form:

$$\rho(-\alpha(X_t)) \leq \delta \Leftrightarrow \delta - \rho(-\alpha(X_t)) \geq 0 \quad (7)$$

$$\rho(\alpha(X_t)) \leq \delta \Leftrightarrow \delta - \rho(\alpha(X_t)) \geq 0. \quad (8)$$

### B. Risk-tightened predicates

Having found atomic predicate risk constraints (7) and (8), we now turn to reformulating the stochastic STL problem into a deterministic one assuming affine predicates<sup>1</sup>. Consider the following rearrangement of (2):

$$\begin{aligned} X_t &= \overbrace{A_{t-1}A_{t-2} \dots A_0}^{:=G_{t-1}} x_0 \\ &+ \overbrace{[A_{t-1} \dots A_1 B_0 \quad \dots \quad A_{t-1} B_{t-2} \quad B_{t-1}]}^{:=H_{t-1}} [u_0^T \dots u_{t-1}^T]^T \\ &+ \overbrace{[A_{t-1} \dots A_1 \quad \dots \quad A_{t-1} \quad I]}^{:=L_{t-1}} [W_0^T \dots W_{t-1}^T]^T \\ &= G_{t-1}x_0 + H_{t-1}u_{0:t-1} + L_{t-1}W_{0:t-1}. \end{aligned} \quad (9)$$

Any random variable can be written as a sum of its expectation and a zero-mean random variable. Let  $\bar{W}_{0:t-1} := \mathbb{E}[W_{0:t-1}]$  ( $\bar{W}_t := \mathbb{E}[W_t]$ ). We thus have:

$$X_t = \bar{x}_t + \hat{X}_t \quad (10)$$

$$\bar{x}_t = \mathbb{E}[X_t] = G_{t-1}x_0 + H_{t-1}u_{0:t-1} + L_{t-1}\bar{W}_{0:t-1} \quad (11)$$

$$\hat{X}_t = L_{t-1}(W_{0:t-1} - \bar{W}_{0:t-1}) \quad (12)$$

It is easy to see that (11) results in the deterministic system:

$$\bar{x}_t = A_{t-1}\bar{x}_{t-1} + B_{t-1}\bar{u}_{t-1} + \bar{W}_{t-1}. \quad (13)$$

Using (10) in (7) and (8), assuming affine predicates  $\alpha(X_t) = a^T X_t + b$ , and using *translation invariance* of  $\rho$  we get:

$$\delta - \rho(\pm\alpha(X_t)) \geq 0 \Leftrightarrow \delta - \rho(\pm a^T \hat{X}_t) \mp \alpha(\bar{x}_t) \geq 0 \quad (14)$$

<sup>1</sup>Note that the use of affine predicates only is not particularly restrictive since most Mixed-Integer Linear Programming (MILP) tools, widely used with STL, only allow affine predicates.

Notice that  $\bar{\delta}_{\rho*} : \mathbb{R} \times \mathbf{X} \rightarrow \mathbb{R} \cup \{\pm\infty\}$  is deterministic ( $* \in \{+, -\}$ ). Thus, (14) is affine in  $\bar{x}_t$  and can be used as an *affine deterministic predicate function*.

**Definition 2.** *Risk-Tightened Predicates.* Given risk metric  $\rho$ , risk bound  $\delta \in \mathbb{R}$ , affine predicate function  $\alpha$  and random variable  $X_t$  per (10), we define risk-tightened affine predicates as:

$$\begin{aligned} \bar{\pi} &:= \{\bar{\alpha}_{\rho*}(\delta, X_t) \geq 0\} \\ &:= \begin{cases} \{\bar{\delta}_{\rho-}(\delta, \hat{X}_t) + \alpha(\bar{x}_t) \geq 0\} & \text{if } * = - \\ \{\bar{\delta}_{\rho+}(\delta, \hat{X}_t) - \alpha(\bar{x}_t) \geq 0\} & \text{if } * = + \end{cases} \end{aligned} \quad (15)$$

We now present a corollary that establishes the form of the deterministic STL constraint and the corresponding deterministic system.

**Corollary 1.** *Consider STL formula  $\varphi$  given by (1). The risk constraint  $\mathcal{R}((\Xi_N, t) \# \varphi) \leq \delta$  on (2) with affine predicates is equivalent to a deterministic STL constraint  $\bar{\varphi}$  on the mean trajectory  $\bar{\xi}_N = \bar{x}_1 \dots \bar{x}_{t+N}$  of (2) given by (13) with affine predicates per Definition 2.*

*Proof.* The proof follows from the reformulation of (2) as (9) and using (10) in the affine predicate  $\alpha(X_t)$ . This results in the deterministic system (13) and  $\rho(\pm\alpha(X_t)) \leq \delta$  iff  $\bar{\alpha}_{\rho\pm}(\delta, X_t) \geq 0$ .  $\bar{\varphi}$  follows from Definitions 1 and 2.  $\square$

**Example 1.** Consider the atomic predicates  $\pi_1, \pi_2, \pi_3$  corresponding to  $\{\alpha_1(x) \geq 0\}, \{\alpha_2(x) \geq 0\}, \{\alpha_3(x) \geq 0\}$  respectively (affine predicate functions) and the formula  $\varphi = \pi_1 \mathcal{U}_{[3,5]}(\pi_2 \wedge \neg\pi_3)$ . The risk of violating  $\varphi$  with  $\delta$  risk bound is  $\mathcal{R}((\Xi_5, 0) \# \varphi) \leq \delta$ . Theorem 1 and Definition 1 yield

$$\begin{aligned} \mathcal{R}((\Xi_5, 0) \# \pi_1 \mathcal{U}_{[3,5]}(\pi_2 \wedge \neg\pi_3)) &\leq \delta \\ \Leftrightarrow \exists t' \in [3, 5] \mathcal{R}((\Xi_5, t') \# \pi_2) &\leq \delta \wedge \mathcal{R}((\Xi_5, t') \# \neg\pi_3) \leq \delta \\ &\wedge \forall t'' \in [0, t'] \mathcal{R}((\Xi_5, t'') \# \pi_1) \leq \delta \\ \Leftrightarrow \exists t' \in [3, 5] \rho(-\alpha_2(X_{t'})) &\leq \delta \wedge \rho(\alpha_3(X_{t'})) \leq \delta \\ &\wedge \forall t'' \in [0, t'] \rho(-\alpha_1(X_{t''})) \leq \delta. \end{aligned}$$

Using Definition 2 and letting  $\bar{\pi}_1 = \{\bar{\alpha}_{1\rho-} \geq 0\}$ ,  $\bar{\pi}_2 = \{\bar{\alpha}_{2\rho-} \geq 0\}$ , and  $\bar{\pi}_3 = \{\bar{\alpha}_{3\rho+} \geq 0\}$  result in

$$\begin{aligned} \mathcal{R}((\Xi_5, 0) \# \pi_1 \mathcal{U}_{[3,5]}(\pi_2 \wedge \neg\pi_3)) &\leq \delta \\ \Leftrightarrow \exists t' \in [3, 5] \bar{\alpha}_{2\rho-}(\delta, X_{t'}) &\geq 0 \wedge \bar{\alpha}_{3\rho+}(\delta, X_{t'}) \geq 0 \\ &\wedge \forall t'' \in [0, t'] \bar{\alpha}_{1\rho-}(\delta, X_{t''}) \geq 0 \\ \Leftrightarrow \exists t' \in [3, 5] (\bar{\xi}_5, t') \models \bar{\pi}_2 \wedge &(\bar{\xi}_5, t') \models \bar{\pi}_3 \\ &\wedge \forall t'' \in [0, t'] (\bar{\xi}_5, t'') \models \bar{\pi}_1 \\ \Leftrightarrow (\bar{\xi}_5, 0) \models \bar{\pi}_1 \mathcal{U}_{[3,5]}(\bar{\pi}_2 \wedge \bar{\pi}_3) &\Rightarrow \bar{\varphi} := \bar{\pi}_1 \mathcal{U}_{[3,5]}(\bar{\pi}_2 \wedge \bar{\pi}_3). \end{aligned}$$

*C. Explicit reformulation for affine predicates and Distributionally Robust Value at Risk*

In Section V-A we obtained constraints on atomic predicates of the form:  $\rho(-\alpha(X_t)) \leq \delta$  or  $\rho(\alpha(X_t)) \leq \delta$ . In Section V-B we reformulated the system, derived risk-tightened atomic predicate constraints with affine predicates ( $\alpha(x) = a^T x + b$ ), and presented the deterministic system. In this section we turn to evaluating the tightened constraints for a particular choice

of the risk metric  $\rho$ : the Distributionally Robust Value at Risk (DR-VaR). Under DR-VaR with ambiguity set  $\mathcal{P} = \mathcal{P}_t^W$  from Section IV, the atomic predicate risk constraint, with  $\delta$  constrained to (0, 1), can be rewritten as:

$$\begin{aligned} \rho(\alpha(X_t)) \leq \delta &\Leftrightarrow \sup_{\mathbb{P} \in \mathcal{P}} \mathbb{P}[a^T X_t + b \geq 0] \leq \delta \\ &\Leftrightarrow \inf_{\mathbb{P} \in \mathcal{P}} \mathbb{P}[a^T X_t + b \leq 0] \geq 1 - \delta \end{aligned} \quad (16)$$

$$\begin{aligned} \rho(-\alpha(X_t)) \leq \delta &\Leftrightarrow \sup_{\mathbb{P} \in \mathcal{P}} \mathbb{P}[-a^T X_t - b \geq 0] \leq \delta \\ &\Leftrightarrow \inf_{\mathbb{P} \in \mathcal{P}} \mathbb{P}[-a^T X_t - b \leq 0] \geq 1 - \delta \end{aligned} \quad (17)$$

We now present Lemma 1 which explicitly presents the deterministic constraints equivalent to (16) and (17).

**Lemma 1.** *The DR-VaR constraints (16) and (17) are equivalent to the deterministic affine inequality constraints*

$$a^T \bar{x}_t + b + \Delta \|\Sigma_{W_{0:t-1}}^{1/2} L_{t-1}^T a\|_2 \leq 0, \quad (18)$$

$$a^T \bar{x}_t + b - \Delta \|\Sigma_{W_{0:t-1}}^{1/2} L_{t-1}^T a\|_2 \geq 0 \quad (19)$$

respectively, where  $\Delta = \sqrt{(1-\delta)/\delta}$ ,  $\Sigma_{W_{0:t-1}} = \mathbb{E}[(W_{0:t-1} - \bar{W}_{0:t-1})(W_{0:t-1} - \bar{W}_{0:t-1})^T]$ ,  $\bar{W}_{0:t-1} = \mathbb{E}[W_{0:t-1}]$ ,  $\bar{x}_t$  is the deterministic state given by (13) with  $\bar{x}_0 = x_0$ .

*Proof.* The proof follows by using (13) in (16) and (17) and applying [29, Theorem 3.1] to both.  $\square$

## VI. NUMERICAL EXPERIMENTS

Consider  $X := (X_{x1}, X_{y1}, X_{x2}, X_{y2})$  to be a four dimensional random variable and the following formula:

$$\varphi := \square_{[0,3]}(\varphi_{a1} \wedge \varphi_{a2}) \wedge \diamond_{[2,3]} \varphi_r \wedge \square_{[1,3]} \varphi_f \quad (20)$$

where  $\varphi_{a1} := \pi_1 \vee \pi_2 \vee \pi_3 \vee \pi_4$ ,  $\varphi_{a2} := \pi_5 \vee \pi_6 \vee \pi_7 \vee \pi_8$ ,  $\varphi_r := \pi_9 \wedge \pi_{10} \wedge \pi_{11} \wedge \pi_{12}$ , and  $\varphi_f := \pi_{13} \wedge \pi_{14} \wedge \pi_{15} \wedge \pi_{16}$  encode two agents  $ag_1 := (X_{x1}, X_{y1})$  and  $ag_2 := (X_{x2}, X_{y2})$  avoiding a square set of side length 1 centered at  $P_a := (0, 0)$  and agent  $ag_1$  reaching a square set of side length 1 centered at  $P_r := (2, 0)$  while the two retain a maximum infinity norm distance  $d := 1$  of each other. We have:  $\alpha_1(X) := -X_{x1} - 0.5$ ,  $\alpha_2(X) := X_{x1} - 0.5$ ,  $\alpha_3(X) := -X_{y1} - 0.5$ ,  $\alpha_4(X) := X_{y1} - 0.5$ ,  $\alpha_5 - \alpha_8$  are similar,  $\alpha_9(X) := X_{x1} - 1.5$ ,  $\alpha_{10}(X) := 2.5 - X_{x1}$ ,  $\alpha_{11}(X) := 0.5 + X_{y1}$ ,  $\alpha_{12}(X) := 0.5 - X_{y1}$ ,  $\alpha_{13}(X) := 1 - X_{x1} + X_{x2}$ ,  $\alpha_{14}(X) := 1 + X_{x1} - X_{x2}$ ,  $\alpha_{15}(X) := 1 - X_{y1} + X_{y2}$ ,  $\alpha_{16}(X) := 1 + X_{y1} - X_{y2}$ .

We use the BluSTL [30] package in its closed-loop deterministic setting for our results. The package uses YALMIP [31] and reformulates an STL specification into a Mixed-Integer Linear Program (MILP), and then solves an MPC problem. Each agent is a double integrator with four states (position and velocity along two directions), two control inputs (along velocity directions), and four additive disturbances on the states. The disturbances are sampled from a 0-mean 3 degree of freedom t-distribution scaled to 0.005-variance and applied to velocity states only. BluSTL converts the continuous time system to discrete-time with 0.1sec system steps and 0.2sec controller steps. We perform two types of simulations. Type 1: The controller ignores the disturbance in the dynamics

and uses deterministic MPC with the nominal model, but is evaluated with the disturbance in closed-loop. Type 2: We explicitly incorporate the uncertainty using our proposed STL risk analysis, transform the specification into a deterministic risk-tightened STL formula, compute the control, and then evaluate in closed-loop with the disturbance. We use  $\delta_1 = 0.1$  risk bound for predicates involving the obstacle avoidance and  $\delta_2 = 0.5$  risk bound for predicates involving the goal region. This puts more emphasis on obstacle avoidance. We also saturate the tightening parameter  $\Delta \|\Sigma_{W_{0:T-1}}^{1/2} L_{T-1}^T a\|_2$  after 1 second to limit the effect of increasing prediction uncertainty. In both cases the objective  $J$  minimizes the sum of the 1-norm of the control inputs and the distance to the goal. We run 100 simulations for each type and present them in Fig. 1.




Fig. 1. Type 1 (left) and type 2 (right) simulations of 100 sample trajectories for the task (20). Agent 1 (solid blue) and agent 2 (dashed orange) start from the center of the two blue circles, must avoid the obstacle (red square with cross) and agent 1 must reach the goal (green square).

Type 1 trajectories (left) get very close and occasionally collide with the obstacle. The disturbance causes 84 of the 100 simulations to fail to reach the goal. Type 2 (right) trajectories, however, satisfy the risk bounds and remain sufficiently far from the obstacle, and 98 of the 100 reach the goal.

## VII. CONCLUSION AND FUTURE WORK

We presented a general framework for risk-based STL specifications for stochastic systems using axiomatic risk theory. We are exploring several extensions and variations in ongoing and future work, including explicit reformulations for various risk measures and ambiguity sets, non-affine predicates, non-linear dynamics, infinite-horizon persistent tasks, and alternative risk semantics.

## REFERENCES

- [1] C. Baier and J.-P. Katoen, *Principles of Model Checking*, 1st ed. Cambridge, MA: The MIT Press, 2008.
- [2] H. Kress-Gazit, G. E. Fainekos, and G. J. Pappas, “Temporal-logic-based reactive mission and motion planning,” *IEEE Trans. Robot.*, vol. 25, no. 6, pp. 1370–1381, 2009.
- [3] M. Guo and D. V. Dimarogonas, “Multi-agent plan reconfiguration under local LTL specifications,” *Int. Journal Robot. Research*, vol. 34, no. 2, pp. 218–235, 2015.
- [4] M. Kloetzer and C. Belta, “A fully automated framework for control of linear systems from temporal logic specifications,” *IEEE Trans. Autom. Control*, vol. 53, no. 1, pp. 287–297, 2008.
- [5] O. Maler and D. Nickovic, “Monitoring temporal properties of continuous signals,” in *Proc. Int. Conf. FORMATS FTRIFT*, Grenoble, France, Sep. 2004, pp. 152–166.
- [6] G. E. Fainekos and G. J. Pappas, “Robustness of temporal logic specifications for continuous-time signals,” *Theor. Comput. Sci.*, vol. 410, no. 42, pp. 4262–4291, 2009.
- [7] A. Donzé and O. Maler, “Robust satisfaction of temporal logic over real-valued signals,” in *Proc. Int. Conf. FORMATS*, Klosterneuburg, Austria, Sep. 2010, pp. 92–106.
- [8] N. Mehdipour, C.-I. Vasile, and C. Belta, “Average-based robustness for continuous-time signal temporal logic,” in *Proc. Conf. Decis. Control*, Nice, France, Dec. 2019, pp. 5312–5317.
- [9] L. Lindemann and D. V. Dimarogonas, “Robust control for signal temporal logic specifications using discrete average space robustness,” *Automatica*, vol. 101, pp. 377–387, 2019.
- [10] V. Raman, A. Donzé, M. Maasoumy, R. M. Murray, A. Sangiovanni-Vincentelli, and S. A. Seshia, “Model predictive control with signal temporal logic specifications,” in *Proc. Conf. Decis. Control*, Los Angeles, CA, Dec. 2014, pp. 81–87.
- [11] Y. Pant *et al.*, “Fly-by-logic: control of multi-drone fleets with temporal logic objectives,” in *Proc. Int. Conf. Cyber-Physical Syst.*, Porto, Portugal, April 2018, pp. 186–197.
- [12] L. Lindemann and D. V. Dimarogonas, “Efficient automata-based planning and control under spatio-temporal logic specifications,” in *Proc. Am. Control Conf. (accepted, available under arXiv preprint arXiv:1909.11159)*, Denver, CO, July 2020.
- [13] —, “Control barrier functions for signal temporal logic tasks,” *IEEE control systems letters*, vol. 3, no. 1, pp. 96–101, 2018.
- [14] K. Garg and D. Panagou, “Control-lyapunov and control-barrier functions based quadratic program for spatio-temporal specifications,” in *Proc. Conf. Decis. Control*, Nice, France, Dec. 2019, pp. 1422–1429.
- [15] M. Srinivasan, S. Coogan, and M. Egerstedt, “Control of multi-agent systems with finite time control barrier certificates and temporal logic,” in *Proc. Conf. Decis. Control*, 2018, pp. 1991–1996.
- [16] P. Varnai and D. V. Dimarogonas, “Prescribed performance control guided policy improvement for satisfying signal temporal logic tasks,” in *Proc. Am. Control Conf.*, Philadelphia, PA, July 2019, pp. 286–291.
- [17] D. Muniraj, K. G. Vamvoudakis, and M. Farhood, “Enforcing signal temporal logic specifications in multi-agent adversarial environments: A deep q-learning approach,” in *Proc. Conf. Decis. Control*, Miami, FL, Dec. 2018, pp. 4141–4146.
- [18] Z. Xu, S. Saha, B. Hu, S. Mishra, and A. A. Julius, “Advisory temporal logic inference and controller design for semiautonomous robots,” *IEEE Transactions on Automation Science and Engineering*, vol. 16, no. 1, pp. 459–477, 2018.
- [19] S. Sadraddini and C. Belta, “Robust temporal logic model predictive control,” in *Proc. Conf. Comm., Control, Comp.*, Monticello, IL, Sep. 2015, pp. 772–779.
- [20] V. Raman, A. Donzé, D. Sadigh, R. M. Murray, and S. A. Seshia, “Reactive synthesis from signal temporal logic specifications,” in *Proc. Int. Conf. Hybrid Syst.: Comp. Control*, 2015, pp. 239–248.
- [21] D. Sadigh and A. Kapoor, “Safe control under uncertainty with probabilistic signal temporal logic,” in *Proc. Robot.: Science Syst. XII*, June 2016.
- [22] S. S. Farahani, R. Majumdar, V. S. Prabhu, and S. Soudjani, “Shrinking horizon model predictive control with signal temporal logic constraints under stochastic disturbances,” *IEEE Trans. Autom. Control*, vol. 64, no. 8, pp. 3324–3331, 2018.
- [23] P. Artzner, F. Delbaen, J.-M. Eber, and D. Heath, “Coherent measures of risk,” *Mathematical finance*, vol. 9, no. 3, pp. 203–228, 1999.
- [24] A. Majumdar and M. Pavone, “How should a robot assess risk? towards an axiomatic theory of risk in robotics,” in *Robotics Research*. Springer, 2020, pp. 75–84.
- [25] S. Samuelson and I. Yang, “Safety-aware optimal control of stochastic systems using conditional value-at-risk,” in *Proc. Am. Control Conf.*, Milwaukee, WI, June 2018, pp. 6285–6290.
- [26] S. Singh, Y. Chow, A. Majumdar, and M. Pavone, “A framework for time-consistent, risk-sensitive model predictive control: Theory and algorithms,” *IEEE Trans. Autom. Control*, vol. 64, no. 7, pp. 2905–2912, 2018.
- [27] W. Wiesemann, D. Kuhn, and M. Sim, “Distributionally robust convex optimization,” *Op. Research*, vol. 62, no. 6, pp. 1358–1376, 2014.
- [28] S. Zymler, D. Kuhn, and B. Rustem, “Worst-case value at risk of nonlinear portfolios,” *Management Science*, vol. 59, no. 1, pp. 172–188, 2013.
- [29] G. C. Calafiore and L. El Ghaoui, “On distributionally robust chance-constrained linear programs,” *Journal Opt. Theory App.*, vol. 130, no. 1, pp. 1–22, 2006.
- [30] A. Donzé, V. Raman, G. Frehse, and M. Althoff, “Blustl: Controller synthesis from signal temporal logic specifications,” in *ARCH@ CPSWeek*, 2015, pp. 160–168.
- [31] J. Lofberg, “Yalmip: A toolbox for modeling and optimization in matlab,” in *2004 IEEE international conference on robotics and automation (IEEE Cat. No. 04CH37508)*. IEEE, 2004, pp. 284–289.