

Convexity and monotonicity in nonlinear optimal control under uncertainty

Kevin J. Kircher, *Member, IEEE*, and K. Max Zhang

Abstract—We consider the problem of finite-horizon optimal control design under uncertainty for imperfectly observed discrete-time systems with convex costs and constraints. It is known that this problem can be cast as an infinite-dimensional convex program when the dynamics and measurements are linear, uncertainty is additive, and the risks associated with constraint violations and excessive costs are measured in expectation or in the worst case. In this paper, we extend this result to systems with convex or concave dynamics, nonlinear measurements, more general uncertainty structures and other coherent risk measures. In this setting, the optimal control problem can be cast as an infinite-dimensional convex program if (1) the costs, constraints and dynamics satisfy certain monotonicity properties, and (2) the measured outputs can be reversibly ‘purified’ of the influence of the control inputs through Q - or Youla-parameterization. The practical value of this result is that the finite-dimensional subproblems arising in a variety of suboptimal control methods, notably including model predictive control and the Q -design procedure, are also convex for this class of nonlinear systems. Subproblems can therefore be solved to global optimality using convenient modeling software and efficient, reliable solvers. We illustrate these ideas in a numerical example.

Index Terms—Nonlinear systems, optimal control, convex optimization, model predictive control, Q -design, Youla parameterization, scenario optimization, sample-average approximation.

I. INTRODUCTION

We consider the problem of operating a system that evolves in discrete time steps over a finite horizon under uncertainty. Our goal is to design an output feedback control policy that minimizes a convex cost while satisfying convex constraints. This problem has been studied in the stochastic setting, where uncertainty is modeled probabilistically, the expected cost is minimized, and constraints are enforced in expectation or in probability. [1]–[6] It has also been studied in the robust setting, where uncertainty is assumed to come from a given set, the worst-case cost is minimized, and constraints are enforced for all possible realizations of the uncertain influences on the system. [7]–[11] We accommodate both approaches in this paper, as well as others, adopting the flexible view of risk developed in [12]–[17].

Optimal control under uncertainty is a hard problem, even when the costs and constraints are convex. This is due in part to the fact that the optimization variable is a *policy*, a collection of functions that map measured system outputs into control inputs at each time step. In the stochastic setting, optimal policies can, in principle, be found analytically using stochastic dynamic programming. [18] In practice, analytical solution is

typically limited to systems of very low dimension. Notable exceptions are linear systems with additive uncertainty, no constraints, and costs that are either quadratic [19], [20] or exponential-of-quadratic [21], [22].

When analytical solution is impractical, various methods can generate suboptimal solutions that often perform well. Some examples are classical linear feedback control design, reinforcement learning [23]–[26] and approximate dynamic programming [4], [27]–[29], approximation methods for multistage stochastic programming [6], [30]–[33] and robust optimization with recourse [33]–[36], the Q -design procedure [37]–[42], and model predictive control (MPC) in its certainty-equivalent [43]–[45], stochastic [46]–[52], and robust [53]–[57] variants. When perfect state information is not available, control methods may be paired with a state estimator such as a linear [58], [59], extended [60]–[62] or unscented [63], [64] Kalman filter or a particle filter [65], [66].

These methods vary widely in their scope, scalability and performance. A common theme, however, is that they tend to work best for linear systems. This is due in part to the fact that for linear systems, (an equivalent transformation of) the optimal control problem is convex. [41] Suboptimal control methods often involve numerically solving optimization subproblems generated by the original optimal control problem. The convexity of the original problem typically carries over to the subproblems, allowing them to be efficiently and reliably solved to global optimality. When subproblems are nonconvex, however, guarantees of global optimality are generally unavailable, and solvers and initial guesses may need to be carefully tailored to the applications at hand.

In [41], Skaf and Boyd demonstrate that for linear systems with additive uncertainty, the control design problem can be transformed to an infinite-dimensional convex program. Their method hinges on a change of variables related to the Q - or Youla-parameterization [67]–[69] and to purified output feedback control [2], [10], [34]. When perfect state information is available, this change of variables parameterizes state feedback policies by equivalent disturbance feedback policies. Similar arguments have justified the use of (typically affine) disturbance feedback policies in robust and stochastic MPC of perfectly-observed linear systems. [42], [47], [70]–[74]

For nonlinear systems, the optimal control problem is widely understood to be nonconvex due to nonlinear equality constraints introduced by nonlinear dynamics. Equality constraints can be eliminated, however, by iteratively applying the dynamics to express the state trajectory in terms of the control and exogenous input trajectories. In [75], Rantzer

The authors are with the Sibley School of Mechanical and Aerospace Engineering, Cornell University, Ithaca, NY 14853. {kjk82, kz33}@cornell.edu

and Bernhardsson observe that in convex-monotone systems, where the dynamics are convex and nondecreasing in the state and control input, the state trajectory is *convex* in the control input trajectory. In [76], Schmitt *et al.* generalize this result to convex-state-monotone systems, where the dynamics need not be monotone in the control input.

An immediate consequence of the observations in [75], [76] is that for some nonlinear systems, *open-loop* optimal control (where the decision variable is a fixed control trajectory, rather than a feedback policy) is a convex optimization problem. This holds for convex-state-monotone systems in particular, provided the cost and constraints are nondecreasing in the states. This raises two further questions: Are there other nonlinear systems for which open-loop optimal control is convex? What about *closed-loop* optimal control, where we optimize over policies?

This discussion motivates the definition of a *convex system* as one for which open-loop optimal control is a convex optimization problem. After setting the stage in §II, we establish three results for convex systems in this paper.

- 1) **Characterization (§III):** Systems with mixed convex and linear dynamics are convex systems, provided (a) any linear dynamics are independent of states with nonlinear dynamics, and (b) the cost, constraints and nonlinear dynamics are nondecreasing in the states with nonlinear dynamics. Concave dynamics can also be accommodated.
- 2) **Convex closed-loop design (§IV):** If the measured outputs can be reversibly ‘purified’ of the influence of the control inputs, then the closed-loop optimal control problem for convex systems can be transformed to an equivalent infinite-dimensional convex program. The transformation involves changing variables to policies in the purified outputs via Q - or Youla-parameterization.
- 3) **Approximate solution (§V–VI):** The finite-dimensional subproblems arising in a variety of suboptimal control methods, notably including MPC and the Q -design procedure, are convex for convex systems. Subproblems can therefore be solved to global optimality using convenient modeling software and efficient, reliable solvers. We illustrate this in a numerical example.

II. PROBLEM STATEMENT

A. System

We consider a system to be operated over a finite discrete time span $t = 0, \dots, T$. We can influence the system through the control inputs u_0, \dots, u_{T-1} . The system is also influenced by exogenous inputs $\delta_0, \dots, \delta_T$, which are generally uncertain.¹ The exogenous inputs may include process disturbances, sensor noise, initial states, uncertain model parameters, prices, command or reference signals, *etc.* The exogenous inputs could come from bounded or unbounded sets, and need not be independent or identically distributed over time.

¹We view the exogenous input δ_t as a random vector defined on an underlying probability space. We do not assume that the joint distribution of the exogenous input trajectory is known, but we do assume this distribution is independent of the control input trajectory.

The control and exogenous inputs determine the states x_0, \dots, x_T through the system dynamics:

$$\begin{aligned} x_0 &= f_0(\delta_0) \\ x_t &= f_t(x_{t-1}, u_{t-1}, \delta_t), \quad t = 1, \dots, T. \end{aligned} \quad (1)$$

We observe the system through the measured outputs

$$\begin{aligned} y_0 &= g_0(x_0, \delta_0) \\ y_t &= g_t(x_t, u_{t-1}, \delta_t), \quad t = 1, \dots, T-1. \end{aligned} \quad (2)$$

We assume that the dynamics mappings f_0, \dots, f_T and measurement mappings g_0, \dots, g_{T-1} are known. At each time t , the controller receives the measured output y_t . It decides the control input

$$u_t = \pi_t(y_0, \dots, y_t)$$

by evaluating an output feedback control law π_t . The control policy $\pi = (\pi_0, \dots, \pi_{T-1})$ is designed in advance; this design problem is the subject of this paper.

To simplify specification of the policy design problem, we work with the input, state and output trajectories. By iteratively applying the dynamics, the state trajectory can be expressed in terms of the control and exogenous input trajectories as

$$x = \begin{bmatrix} x_0 \\ x_1 \\ \vdots \\ x_T \end{bmatrix} = \begin{bmatrix} \phi_0(\delta_0) \\ \phi_1(u_0, \delta_{0:1}) \\ \vdots \\ \phi_T(u_{0:T-1}, \delta_{0:T}) \end{bmatrix} = \phi(u, \delta).$$

Here the subscript $t_1 : t_2$ denotes a trajectory from time t_1 to t_2 . For example, $\delta_{0:T} = (\delta_0, \dots, \delta_T) = \delta$ and $u_{0:T-1} = (u_0, \dots, u_{T-1}) = u$. The input-state mappings ϕ_t are defined by the recursion

$$\begin{aligned} \phi_0(\delta_0) &= f_0(\delta_0) \\ \phi_1(u_0, \delta_{0:1}) &= f_1(\phi_0(\delta_0), u_0, \delta_1) \\ \phi_t(u_{0:t-1}, \delta_{0:t}) &= f_t(\phi_{t-1}(u_{0:t-2}, \delta_{0:t-1}), u_{t-1}, \delta_t), \\ & \quad t = 2, \dots, T. \end{aligned}$$

Similarly, the measured output trajectory can be written as

$$y = \begin{bmatrix} y_0 \\ y_1 \\ \vdots \\ y_{T-1} \end{bmatrix} = \begin{bmatrix} \psi_0(\delta_0) \\ \psi_1(u_0, \delta_{0:1}) \\ \vdots \\ \psi_{T-1}(u_{0:T-2}, \delta_{0:T-1}) \end{bmatrix} = \psi(u, \delta),$$

where the input-output mappings ψ_t are defined recursively by

$$\begin{aligned} \psi_0(\delta_0) &= g_0(f_0(\delta_0), \delta_0) \\ \psi_t(u_{0:t-1}, \delta_{0:t}) &= g_t(\phi_t(u_{0:t-1}, \delta_{0:t}), u_{t-1}, \delta_t), \\ & \quad t = 1, \dots, T-1. \end{aligned}$$

We also write the control input trajectory as

$$u = \pi(y) = \begin{bmatrix} \pi_0(y_0) \\ \pi_1(y_{0:1}) \\ \vdots \\ \pi_{T-1}(y_{0:T-1}) \end{bmatrix}.$$

Figure 1 illustrates the system.

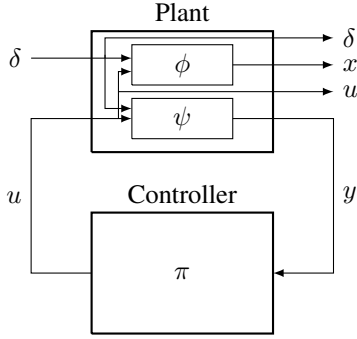


Fig. 1. The closed-loop system is driven by the exogenous input δ . The control input u is determined by the measured output y through the control policy π . Performance is evaluated based on δ , u and the state x .

B. Cost, constraints and risk measures

We are interested in designing the output feedback control policy $\pi = (\pi_0, \dots, \pi_{T-1})$. The policy must be *causal*, meaning each control law can depend on outputs from the past or present, but not the future. We denote the set of causal output feedback policies by Π . Our goal is to find a $\pi \in \Pi$ that minimizes a cost

$$c_0(x, u, \delta) \quad (3)$$

while satisfying constraints

$$c_j(x, u, \delta) \leq 0, \quad j = 1, \dots, J. \quad (4)$$

We assume that each scalar-valued function c_0, \dots, c_J is convex in (x, u) for all δ .

Because the exogenous inputs are uncertain, the goals of minimizing cost and satisfying constraints are ambiguous. Should they be accomplished on average with respect to the distribution of δ , for all possible realizations of δ , or somewhere in between? Resolving these ambiguities amounts to choosing measures of risk.² A *risk measure* \mathcal{R} is a functional that maps uncertain scalars into deterministic scalars (or possibly ∞). It quantifies the risk associated with positive realizations of its uncertain argument. In the language of [12], [13], the risk measure \mathcal{R}_j should be associated with constraint j if the designer views the risk of constraint violation as *acceptable* whenever

$$\mathcal{R}_j c_j(x, u, \delta) \leq 0.$$

The risks associated with excessive costs can be treated similarly to the risks associated with constraint violations, because minimizing $c_0(x, u, \delta)$ is equivalent to minimizing a deterministic scalar τ subject to the constraint $c_0(x, u, \delta) - \tau \leq 0$.

Table I describes five risk measures commonly used in optimization under uncertainty. The expected value is risk-neutral: it weighs constraint violations and slacks equally, modeling the risk of constraint violation as acceptable if slacks balance out violations on average. The worst-case value, by contrast, is maximally risk-averse: it models any possibility

²For more on measuring risk, a rich subject that has received much recent attention, we refer the reader to [12]–[17].

TABLE I
EXAMPLE RISK MEASURES FOR A RANDOM VARIABLE X .

Risk measure	Definition of $\mathcal{R}X$	Risk acceptable if...
Expected value	$\mathbf{E} X$	$\bar{X} \leq 0$ on average
Predicted value	prediction \hat{X} of X	$\hat{X} \leq 0$
Worst-case value	$\sup X$	$X \leq 0$ almost surely
Value-at-risk, $\text{VaR}_\beta(X)$	$\inf\{\alpha \mid \mathbf{P}\{X \leq \alpha\} \geq \beta\}$	$\mathbf{P}\{X \leq 0\} \geq \beta$
Conditional value-at-risk	$\inf_\alpha\{\alpha + \mathbf{E}(X - \alpha)_+ / (1 - \beta)\}$	$\mathbf{E}[X \mid X \geq \text{VaR}_\beta(X)] \leq 0$

of constraint violation as unacceptable. The value-at-risk at confidence level $\beta \in [0, 1]$ (VaR_β) is the smallest threshold that, with $100\beta\%$ confidence, will not be exceeded. Typical values of β are 0.9, 0.95 and 0.99. VaR_β generates chance constraints of the form

$$\mathbf{P}\{c_j(x, u, \delta) \leq 0\} \geq \beta.$$

The conditional value-at-risk at confidence level β (CVaR_β) is the conditional expectation, conditioned on the event that VaR_β is exceeded. [14], [15] CVaR_β always upper bounds VaR_β . The CVaR_β of $c_j(x, u, \delta)$ can be constrained by minimizing a scalar α subject to

$$\mathbf{E}(c_j(x, u, \delta) - \alpha)_+ \leq -\alpha(1 - \beta),$$

where $(\cdot)_+$ denotes the positive part, $\max\{0, \cdot\}$. An optimal α^* upper bounds $\text{VaR}_\beta c_j(x, u, \delta)$; this bound is usually tight. [15]

In this paper, we leave the choice of risk measures to the designer. This flexibility makes the results in the following sections applicable in a range of settings, including the robust and stochastic frameworks and others. We do, however, impose one restriction on the risk measures: we assume that they are *coherent* in the sense of [12], [13], [16], [17]. By this, we mean that each risk measure is convex, nondecreasing, lower semicontinuous, and preserves certainty (*i.e.*, $\mathcal{R}(C) = C$ for constant C). We make this assumption primarily because coherent risk measures preserve convexity: if \mathcal{R}_j is coherent and $(x, u) \mapsto c_j(x, u, \delta)$ is convex for all δ , then $(x, u) \mapsto \mathcal{R}_j c_j(x, u, \delta)$ is convex. [16]

All but one of the risk measures mentioned above are coherent. The exception is VaR_β , which is convex only in structured special cases. We note, however, that VaR_β can be approximated in a convexity-preserving fashion using sampling [77]–[82] or conservative upper bounds such as CVaR_β [6], [50].

In summary, we are interested in the following problem of optimal control under uncertainty:

$$\begin{aligned} & \underset{\pi \in \Pi}{\text{minimize}} && \mathcal{R}_0 c_0(x, u, \delta) \\ & \text{subject to} && \mathcal{R}_j c_j(x, u, \delta) \leq 0, \quad j = 1, \dots, J \\ & && x = \phi(u, \delta) \\ & && u = \pi(y) \\ & && y = \psi(u, \delta). \end{aligned} \quad (\text{OC})$$

This problem is intractable in its current form. This is due in part to the fact that the decision variable π is infinite-dimensional, an obstacle that we will address approximately in §V. Problem (OC) is also complicated by the interdependence

between the state, output and control trajectories. We will unravel this interdependence in §IV using a nonlinear change of variables. First, we will consider the case of open-loop control, where the control laws π_t are restricted to be constant.

III. CONVEX SYSTEMS

To understand when closed-loop control design is a convex optimization problem, we begin by building intuition in the simpler context of open-loop control. In this context, a (finite-dimensional) vector of control inputs $\tilde{u} = (\tilde{u}_0, \dots, \tilde{u}_{T-1})$ is decided in advance and implemented without feedback. This amounts to drastically restricting the search space from the set of all causal output feedback policies to the subspace of *constant* policies. With this restriction, the open-loop optimal control problem is to

$$\begin{aligned} & \underset{\tilde{u}}{\text{minimize}} && \mathcal{R}_0 c_0(\phi(\tilde{u}, \delta), \tilde{u}, \delta) \\ & \text{subject to} && \mathcal{R}_j c_j(\phi(\tilde{u}, \delta), \tilde{u}, \delta) \leq 0, \quad j = 1, \dots, J. \end{aligned} \quad (\text{OLOC})$$

This is a finite-dimensional optimization problem in the vector \tilde{u} . Because coherent risk measures preserve convexity [16], Problem (OLOC) is convex if each cost and constraint function c_j , when composed with the input-state mapping ϕ , is convex in u for all δ . This motivates the following definition.

Definition 1 (Convex system). *We call the system with dynamics (1), cost (3) and constraints (4) a convex system if the functions*

$$u \mapsto c_j(\phi(u, \delta), u, \delta), \quad j = 0, \dots, J$$

are convex for all δ .

By definition, open-loop optimal control of convex systems is a convex optimization problem. In §IV, we will show that *closed-loop* optimal control of convex systems is also convex, provided the outputs can be reversibly ‘purified’ of the influence of the control inputs. First, we will characterize a class of convex systems.

A. Characterizing convex systems

To understand when a system is convex, we need to understand how the states and controls propagate through the dynamics (1) into the cost (3) and constraints (4). If the dynamics are linear, then this process is straightforward. The state trajectory x is affine in the control trajectory u for all δ . Convexity is preserved under composition with an affine mapping [83], so for linear systems $u \mapsto c_j(\phi(u, \delta), u, \delta)$ is convex for all δ .

When the system is nonlinear, the process is less straightforward. We do not explore it in general here. Instead, we consider a class of systems with the following dynamics:

$$x_t = \begin{bmatrix} x_t^{\text{aff}} \\ x_t^{\text{cvx}} \end{bmatrix} = \begin{bmatrix} A_t(\delta_t)x_{t-1}^{\text{aff}} + B_t(\delta_t)u_{t-1} + w_t(\delta_t) \\ h_t(x_{t-1}^{\text{aff}}, x_{t-1}^{\text{cvx}}, u_{t-1}, \delta_t) \end{bmatrix}. \quad (5)$$

For this class of systems, convexity can be established using simple composition rules. An important restriction is that the states x_t^{aff} with linear dynamics are independent of the nonlinear states x_t^{cvx} . This ensures that the trajectory $x^{\text{aff}} =$

$(x_0^{\text{aff}}, \dots, x_T^{\text{aff}})$ is affine in the control trajectory. Any states with linear dynamics that depend on nonlinear states are included in x_t^{cvx} .

Theorem 1. *The system with dynamics (5), cost (3) and constraints (4) is a convex system if the following conditions hold for all δ .*

- 1) For $t = 1, \dots, T$, each row of the nonlinear dynamics mapping h_t is
 - a) jointly convex in $(x_{t-1}^{\text{aff}}, x_{t-1}^{\text{cvx}}, u_{t-1})$, and
 - b) nondecreasing in each element of x_{t-1}^{cvx} .
- 2) For $j = 0, \dots, J$, the function

$$(x^{\text{aff}}, x^{\text{cvx}}, u) \mapsto c_j((x^{\text{aff}}, x^{\text{cvx}}), u, \delta)$$

is nondecreasing in each element of $x^{\text{cvx}} = (x_0^{\text{cvx}}, \dots, x_T^{\text{cvx}})$.

Appendix A contains a proof of Theorem 1. The proof hinges on the fact that the composition of a convex nondecreasing function with a convex function is convex. [83]

Three remarks are in order. First, system (5) is a straightforward generalization of the convex-state-monotone system studied in [76], which in turn generalizes the convex-monotone system studied in [75]. A precise name for system (5) would be *convex-nonlinear-state-monotone*, since the monotonicity requirement applies only to the states with nonlinear dynamics. The purpose of this generalization is to absorb the class of linear systems for which optimal control was shown in [41] to be convex. In particular, the linear-quadratic system is not convex-monotone or convex-state-monotone in general, but is a special case of system (5) with linear dynamics, additive uncertainty, no constraints and quadratic cost.

Second, concave dynamics can be included in system (5) after a sign change. For example, we consider a system with

$$\begin{aligned} z_t^{\text{cvx}} &= h_t^{\text{cvx}}(x_{t-1}^{\text{aff}}, z_{t-1}^{\text{cvx}}, z_{t-1}^{\text{ccv}}, u_{t-1}, \delta_t) \\ z_t^{\text{ccv}} &= h_t^{\text{ccv}}(x_{t-1}^{\text{aff}}, z_{t-1}^{\text{cvx}}, z_{t-1}^{\text{ccv}}, u_{t-1}, \delta_t), \end{aligned}$$

where each row of h_t^{cvx} is nondecreasing in z_{t-1}^{cvx} , nonincreasing in z_{t-1}^{ccv} and convex. Similarly, each row of h_t^{ccv} is nondecreasing in z_{t-1}^{cvx} , nonincreasing in z_{t-1}^{cvx} and concave. If we define the nonlinear state as

$$\begin{aligned} x_t^{\text{cvx}} &= \begin{bmatrix} z_t^{\text{cvx}} \\ -z_t^{\text{ccv}} \end{bmatrix} = \begin{bmatrix} h_t^{\text{cvx}}(x_{t-1}^{\text{aff}}, z_{t-1}^{\text{cvx}}, z_{t-1}^{\text{ccv}}, u_{t-1}, \delta_t) \\ -h_t^{\text{ccv}}(x_{t-1}^{\text{aff}}, z_{t-1}^{\text{cvx}}, z_{t-1}^{\text{ccv}}, u_{t-1}, \delta_t) \end{bmatrix} \\ &= h_t(x_{t-1}^{\text{aff}}, x_{t-1}^{\text{cvx}}, u_{t-1}, \delta_t), \end{aligned}$$

then each row of h_t is nondecreasing in x_{t-1}^{cvx} and convex, as required by Theorem 1.

Third, system (5) includes a rich class of nonlinear systems. Provided careful attention is paid to curvature and monotonicity, it can admit exponentials and logarithms, quadratic forms, roots, powers, nonsmooth functions such as maxima, minima and absolute values, and sums and compositions of the above. Many more examples of nonlinear convex and concave functions can be found in §3 of [83]. Systems with nonlinear convex dynamics have arisen in applications ranging from cancer and HIV treatment scheduling [75] to voltage control in power systems [75] to freeway congestion management [76] to energy storage control [84], [85]. The drug treatment model

in [75] is a special case of a more general class of bilinear systems that, through a logarithmic transformation, can be cast as convex systems.

IV. CONVEX CLOSED-LOOP OPTIMAL CONTROL

By contrast to the open-loop optimal control problem (OLOC), the equality constraints in the closed-loop optimal control problem (OC) cannot be easily eliminated. This is due to the complicated interdependence between the state, output and control input trajectories. Under some conditions on the uncertainty structure, however, this interdependence can be disentangled through a nonlinear change of variables related to the Q - or Youla-parameterization [67]–[69] and purified output feedback control [2], [10], [34]. We now establish sufficient conditions for this change of variables to be possible.

A. Purifiability and Q -parameterization

Definition 2 (Purifiable). *We call the system with dynamics (1) and measurements (2) purifiable if for each $t = 1, \dots, T-1$, there exist mappings p_t , q_t and ξ_t such that for all $u_{0:t-1}$ and $\delta_{0:t}$,*

$$\begin{aligned} p_t(\psi_{0:t}(u_{0:t-1}, \delta_{0:t}), u_{0:t-1}) &= \xi_t(\delta_{0:t}) \\ \iff q_t(\xi_{0:t}(\delta_{0:t}), u_{0:t-1}) &= \psi_t(u_{0:t-1}, \delta_{0:t}). \end{aligned}$$

We call $e_t = \xi_t(\delta_{0:t})$ the purified output and p_t the purifier.

We note for future reference that y_0 is a function of δ_0 only:

$$y_0 = g_0(x_0, \delta_0) = g_0(f_0(\delta_0), \delta_0).$$

Without loss of generality, therefore, we define p_0 , q_0 and ξ_0 such that

$$p_0(y_0) = y_0, \quad q_0(e_0) = e_0, \quad \xi_0(\delta_0) = g_0(f_0(\delta_0), \delta_0).$$

Theorem 2. *If the system with dynamics (1) and measurements (2) is purifiable, then there exists a one-to-one correspondence between causal output feedback policies π and causal policies Q in the purified output $e = \xi(\delta)$. Furthermore, given a causal Q in e , the unique corresponding causal π in y can be constructed from the following recursion:*

$$\begin{aligned} \pi_0(y_0) &= Q_0(y_0) \\ \pi_t(y_{0:t}) &= Q_t(p_{0:t}(y_{0:t}, \pi_{0:t-1}(y_{0:t-1}))), \\ t &= 1, \dots, T-1. \end{aligned} \quad (6)$$

Theorem 2 is closely related to the nonlinear discrete-time Youla parameterization presented by Wu and Lall in [69]. We include a proof in Appendix B for completeness. We give some examples of purifiable nonlinear systems in §IV-B.

Theorem 2 establishes that if the system is purifiable, then we can optimize over policies Q in the purified output $e = \xi(\delta)$. We interpret e_t as what remains of the output y_t when the effect of the control inputs u_0, \dots, u_{t-1} has been removed. Given a causal Q , the unique corresponding causal π can be recovered. We denote the set of causal Q by \mathcal{Q} .

Under this change of variables, an equivalent reformulation of the closed-loop optimal control problem (OC) is to

$$\begin{aligned} &\text{minimize}_{Q \in \mathcal{Q}} \mathcal{R}_0 c_0(x, u, \delta) \\ &\text{subject to} \quad \mathcal{R}_j c_j(x, u, \delta) \leq 0, \quad j = 1, \dots, J \\ &\quad x = \phi(u, \delta) \\ &\quad u = Q(e) \\ &\quad e = \xi(\delta). \end{aligned}$$

The equality constraints can now be eliminated, giving another equivalent problem:

$$\begin{aligned} &\text{minimize}_{Q \in \mathcal{Q}} \mathcal{R}_0 c_0(\phi(Q(\xi(\delta)), \delta), Q(\xi(\delta)), \delta) \\ &\text{subject to} \quad \mathcal{R}_j c_j(\phi(Q(\xi(\delta)), \delta), Q(\xi(\delta)), \delta) \leq 0, \\ &\quad j = 1, \dots, J. \end{aligned} \quad (\text{OC-}Q)$$

Problem (OC- Q) is structurally identical to the open-loop optimal control problem (OLOC), except that the optimization is over (infinite-dimensional) purified output feedback policies Q rather than (finite-dimensional) control input trajectories \tilde{u} . It follows that the two problems are convex under the same conditions, namely for convex systems. This observation, together with Theorem 2, gives the following result.

Corollary 3. *For purifiable systems, the output feedback policy design problem (OC) and the purified output feedback policy design problem (OC- Q) are equivalent. For convex systems, Problem (OC- Q) is convex.*

We will see in §V that Corollary 3 establishes the basic tractability of a variety of suboptimal control schemes for nonlinear convex systems. First, we provide a few examples of purifiable systems.

B. Purifiability examples

Purifiability is essentially an invertibility property on the input-output mapping ψ with respect to the exogenous inputs. The requirements for purifiability are (1) at each time step, the influence of the control input history can be removed from the current output, possibly using the output history; and (2) this process must be reversible, in the sense that the current output can be reconstructed from the purified output history and the control input history. To ground this notion, we now provide some concrete examples of purifiable systems. This list is not exhaustive.

a) *Measured exogenous inputs:* If the exogenous inputs are measured exactly ($y_t = \delta_t$), then the system is purifiable with

$$\begin{aligned} e_t &= \delta_t \\ p_t(y_{0:t}, u_{0:t-1}) &= y_t \\ q_t(e_{0:t}, u_{0:t-1}) &= e_t. \end{aligned}$$

b) *Pure estimation:* If the controller can observe the system but not influence it, then the states and outputs can be expressed as

$$\begin{aligned} x_t &= f_t(x_{t-1}, \delta_t) = \phi_t(\delta_{0:t}) \\ y_t &= g_t(x_t, \delta_t) = g_t(\phi_t(\delta_{0:t}), \delta_t) = \psi_t(\delta_{0:t}). \end{aligned}$$

In this case, the system is trivially purifiable with $e_t = y_t$. This implies that various constrained estimation problems can be

put in the form of Problem (OC- Q). For example, we consider the problem of designing a state estimator π such that $u_t = \pi_t(y_{0:t})$ minimizes the mean squared error in estimating x_t , with the prior knowledge that $x_t \succeq 0$ almost surely. This can be put in the form of Problem (OC- Q) by setting $c_0(x, u) = \|x - u\|_2^2$, $\mathcal{R}_0 = \mathbf{E}$, $c_1(u) = \max_{i,t} \{-(u_t)_i\}$ and $\mathcal{R}_1 = \text{sup}$.

c) *Perfect state information, invertible dynamics*: If the states are measured exactly ($y_t = x_t$) and the dynamics are invertible in the exogenous inputs, *i.e.*, there exist mappings f_t^{-1} such that

$$f_t^{-1}(f_t(x_{t-1}, u_{t-1}, \delta_t), x_{t-1}, u_{t-1}) = \delta_t,$$

then the system is purifiable with

$$\begin{aligned} e_t &= \delta_t \\ p_t(y_{0:t}, u_{0:t-1}) &= f_t^{-1}(y_t, y_{t-1}, u_{t-1}) \\ q_t(e_{0:t}, u_{0:t-1}) &= \phi_t(u_{0:t-1}, e_{0:t}). \end{aligned}$$

In the special case of additive disturbances, $x_t = f_t(x_{t-1}, u_{t-1}) + \delta_t$, the purifier reduces to

$$p_t(y_{0:t}, u_{0:t-1}) = y_t - f_t(y_{t-1}, u_{t-1}).$$

d) *Deterministic dynamics, invertible measurements*: If the initial state is measured exactly ($y_0 = x_0$), the dynamics are deterministic,

$$x_t = f_t(x_{t-1}, u_{t-1}) = \phi_t(u_{0:t-1}, x_0),$$

and the measurements are invertible in the exogenous inputs, *i.e.*, there exist mappings g_t^{-1} such that

$$g_t^{-1}(g_t(x_t, u_{t-1}, \delta_t), x_t, u_{t-1}) = \delta_t,$$

then the system is purifiable with

$$\begin{aligned} e_t &= \delta_t \\ p_t(y_{0:t}, u_{0:t-1}) &= g_t^{-1}(y_t, \phi_t(u_{0:t-1}, y_0), u_{t-1}) \\ q_t(e_{0:t}, u_{0:t-1}) &= g_t(\phi_t(u_{0:t-1}, e_0), u_{t-1}, e_t). \end{aligned}$$

In the special case of additive noise, $y_t = g_t(x_t, u_{t-1}) + \delta_t$, the purifier reduces to

$$p_t(y_{0:t}, u_{0:t-1}) = y_t - g_t(\phi_t(u_{0:t-1}, y_0), u_{t-1}).$$

e) *State-affine dynamics and measurements, additive uncertainty*: If the dynamics and measurements have the form

$$\begin{aligned} f_t(x_{t-1}, u_{t-1}, \delta_t) &= A_t x_{t-1} + f_t^u(u_{t-1}) + w_t(\delta_t) \\ g_t(x_t, u_{t-1}, \delta_t) &= C_t x_t + g_t^u(u_{t-1}) + v_t(\delta_t), \end{aligned}$$

then it can be shown that the input-output mappings are additive:

$$y_t = \psi_t(u_{0:t-1}, \delta_{0:t}) = \psi_t^u(u_{0:t-1}) + \psi_t^\delta(\delta_{0:t}).$$

In this case, the system is purifiable with

$$\begin{aligned} e_t &= \psi_t^\delta(\delta_{0:t}) \\ p_t(y_{0:t}, u_{0:t-1}) &= y_t - \psi_t^u(u_{0:t-1}) \\ q_t(e_{0:t}, u_{0:t-1}) &= e_t + \psi_t^u(u_{0:t-1}). \end{aligned}$$

V. APPROXIMATE SOLUTION METHODS

Although Problem (OC- Q) is convex for convex systems, it remains challenging for two reasons. First, the decision variable Q is infinite-dimensional. Second, the risk measures $\mathcal{R}_0, \dots, \mathcal{R}_J$ may be difficult to compute, or even ill-defined if some distributional information is lacking. We discuss methods for addressing the infinite-dimensionality in §V-A and for approximating risk measures in §V-B. The upshot of this discussion is that several suboptimal control methods that perform well for linear systems can be applied to nonlinear convex systems using finite-dimensional convex optimization.

A. Finite-dimensional restrictions

1) *Open-loop model predictive control*: A simple, effective method for overcoming the challenge of infinite dimensionality is open-loop MPC. In this method, at each time step we solve a version of the (finite-dimensional) open-loop optimal control problem (OLOP) over a truncated, receding horizon. This generates a planned control input trajectory. We implement the first step in this plan, the system evolves, we update the state estimate and the process repeats. As discussed in §III, open-loop optimal control is convex for convex systems, even those with nonlinear dynamics.

In one common variant of open-loop MPC, typically called certainty-equivalent MPC, the subproblem at each time step is solved under a single prediction of the disturbance trajectory. [43]–[45] This amounts to measuring risk with the *predicted value* risk measure discussed in §II-B. Other risk measures can be used in open-loop MPC, however, and can significantly improve performance. [46], [48]–[55], [57] Risk measures can be approximated if necessary, as discussed in §V-B.

Open-loop MPC has several advantages. First, it often performs well in practice. Second, the design process is straightforward and intuitive; the primary design decisions are the prediction horizon, the terminal cost and/or constraints with which to augment the subproblems, and the algorithms for prediction and state estimation. A third advantage, and perhaps an underappreciated one, is that open-loop MPC admits very general uncertainty structures, including additive, multiplicative and others. In particular, since it does not require the Youla-type change of variables discussed in §IV-A, open-loop MPC can be applied to systems that are not purifiable.

Open-loop MPC also has several disadvantages. First, its implementation is computationally intensive due to the use of online optimization; this can limit its scalability. (We note, however, that for linearly constrained linear systems with quadratic costs, open-loop MPC can be implemented efficiently using custom solvers that exploit the problem structure. [45]) Second, open-loop MPC yields no closed-form expression for the control policy π ; this complicates analysis of closed-loop stability, robustness, *etc.* Third, the method's performance can suffer somewhat due to the open-loop structure of the optimal control subproblems. This structure ignores the controller's opportunity to respond to future information as it becomes available, *i.e.*, the controller's *recourse*. The Q -design procedure partially addresses these disadvantages.

2) *Q-design*: In open-loop MPC, the policy design procedure is straightforward (choosing a few parameters and sub-routines), but implementation is computationally intensive due to the use of online optimization. In the *Q*-design procedure [37]–[42], [73], by contrast, policy design is computationally intensive, but implementation is extremely efficient. In particular, no online optimization is needed. Another important distinction is that unlike MPC, *Q*-design yields a closed-form control policy. This facilitates analysis and simulation of the closed-loop system.

In *Q*-design, we design a suboptimal causal purified output feedback policy

$$Q_\theta = \sum_{k=1}^K \theta_k Q^{(k)}.$$

Here $Q^{(1)}, \dots, Q^{(K)} \in \mathcal{Q}$ are causal basis policies selected by the designer and $\theta = (\theta_1, \dots, \theta_K) \in \mathbf{R}^K$ is a parameter vector. The parameters are decided by solving a finite-dimensional analogue of Problem (OC-*Q*), with *Q* replaced by Q_θ :

$$\begin{aligned} & \underset{\theta}{\text{minimize}} && \mathcal{R}_0 c_0(\phi(Q_\theta(\xi(\delta)), \delta), Q_\theta(\xi(\delta)), \delta) \\ & \text{subject to} && \mathcal{R}_j c_j(\phi(Q_\theta(\xi(\delta)), \delta), Q_\theta(\xi(\delta)), \delta) \leq 0, \\ & && j = 1, \dots, J \\ & && Q_\theta(\xi(\delta)) = \sum_{k=1}^K \theta_k Q^{(k)}(\xi(\delta)). \end{aligned} \quad (\text{OC-}Q_\theta)$$

Convex constraints on θ can also be added, *e.g.*, to cultivate a particular structure in the control policy. Given Q_θ , the unique corresponding causal output feedback policy can be constructed from the recursion in Theorem 2.

Because Q_θ is affine in θ and convexity is preserved under composition with an affine mapping, the *Q*-design problem (OC- Q_θ) inherits the convexity of Problem (OC-*Q*) for convex systems. It follows that *Q*-design can be applied to convex systems, even nonlinear ones, using finite-dimensional convex optimization. Risk measures in Problem (OC- Q_θ) can be approximated if necessary, as discussed in §V-B.

In principle, *Q*-design can solve the optimal control problem (OC-*Q*) to any degree of accuracy for a sufficiently large basis. In practice, *Q*-design has two disadvantages. First, its performance depends intimately on the choice of basis policies. In [42], Skaf and Boyd explore the natural choice of affine policies in the context of linear systems. While affine *Q*-design performs well in a number of problems, the performance is typically not as good as open-loop MPC. To our knowledge, the general problem of finding good nonlinear basis policies has not yet been solved. A second disadvantage of *Q*-design is that, as discussed in §IV-A, it can be applied only to purifiable systems. If the system is not purifiable, then an implementable output feedback policy may not be recoverable from the *Q*-designed policy.

3) *Closed-loop model predictive control*: The final approximate solution method we discuss, closed-loop MPC, can be viewed as receding horizon *Q*-design. As in open-loop MPC, in closed-loop MPC we solve an optimal control subproblem at each time step over a truncated, receding horizon. The key difference is that the closed-loop MPC subproblems include a recourse model, *i.e.*, a model of the controller's response to future information as it becomes available. More precisely,

the open-loop MPC subproblems are truncated versions of the open-loop optimal control problem (OLOP), while the closed-loop MPC subproblems are truncated versions of the *Q*-design problem (OC- Q_θ). Closed-loop MPC with affine recourse models is developed for linear systems in [42], [47], [50], [51], [56], [57].

Closed-loop MPC addresses the two disadvantages of *Q*-design. Like open-loop MPC, closed-loop MPC can be applied to systems that are not purifiable. To see this, we note that the first step of the recursion in Theorem 2,

$$\pi_0(y_0) = Q_0(y_0),$$

can *always* be implemented, even if later steps cannot. Closed-loop MPC also tends to be less sensitive than *Q*-design to the choice of basis policies; in closed-loop MPC the policy designed at each time step is merely a model of future recourse, while in *Q*-design the policy is the actual source of the implemented control inputs.

Closed-loop MPC can outperform open-loop MPC due to the inclusion of a recourse model. [42] Its design process is similarly straightforward; the only additional step is choosing basis policies. Like open-loop MPC, however, closed-loop MPC involves computationally intensive, online optimization. A closed-form expression for the closed-loop MPC policy is generally not available.

B. Risk measure approximation

The control methods discussed in §V-A all require solving finite-dimensional convex optimization subproblems under uncertainty. These subproblems can be solved exactly in a few special cases. In general, however, we must resort to approximate solution methods. [3], [6], [10]

We now describe one approximate solution method based on sampling. This method can accommodate each of the risk measures discussed in §II-B. We begin by obtaining samples $\delta^{(1)}, \dots, \delta^{(N)}$ from the distribution of the exogenous input trajectory δ , or in the robust setting, from the corresponding uncertainty set. Samples could be obtained, *e.g.*, from historical data or a pseudorandom number generator. Each risk measure \mathcal{R}_j is then replaced by an approximation $\hat{\mathcal{R}}_j^N$ based on the samples $\delta^{(1)}, \dots, \delta^{(N)}$. This results in a convex optimization problem (with random input data) that can be solved using off-the-shelf software. Depending on the underlying risk measures, this sample-based approximate solution scheme is known as *scenario optimization* [77]–[82] or *sample-average approximation* [6], [86]–[88].

Table II shows sample-based approximations for the five risk measures in Table I. The worst-case value and value-at-risk (which generate robust and chance constraints, respectively) can be approximated by maxima over the full sample. Expectation-type risk measures, including conditional value-at-risk, can be approximated using sample averages. Both of these approximations preserve the convexity of the underlying cost or constraint function. Theoretical bounds on violation probabilities for robust and chance constraints are available. [77]–[82] Some convergence results for sample-average approximation can be found in [6], [86]–[88]. Variance

TABLE II
APPROXIMATIONS OF THE RISK MEASURES IN TABLE I BASED ON
SAMPLES $X^{(1)}, \dots, X^{(N)}$ FROM THE DISTRIBUTION OF X .

Risk measure, $\mathcal{R}X$	Sample-based approximation, $\hat{\mathcal{R}}^N(X^{(1)}, \dots, X^{(N)})$
Expected value	$(X^{(1)} + \dots + X^{(N)})/N$
Predicted value	prediction based on $X^{(1)}, \dots, X^{(N)}$
Worst-case value	$\max\{X^{(1)}, \dots, X^{(N)}\}$
Value-at-risk	$\max\{X^{(1)}, \dots, X^{(N)}\}$
Conditional value-at-risk	$\inf_{\alpha} \{\alpha + ((X^{(1)} - \alpha)_+ + \dots + (X^{(N)} - \alpha)_+) / (N(1 - \beta))\}$

reduction methods and hypothesis tests of solution quality can also be applied. [6]

To illustrate this method, we consider the open-loop optimal control problem (OLOC). The sample-based approximation to this problem is to

$$\begin{aligned} & \underset{\tilde{u}}{\text{minimize}} && \hat{\mathcal{R}}_0^N(c_0(\phi(\tilde{u}, \delta^{(1)}), \tilde{u}, \delta^{(1)}), \dots, \\ & && c_0(\phi(\tilde{u}, \delta^{(N)}), \tilde{u}, \delta^{(N)})) \\ & \text{subject to} && \hat{\mathcal{R}}_j^N(c_j(\phi(\tilde{u}, \delta^{(1)}), \tilde{u}, \delta^{(1)}), \dots, \\ & && c_j(\phi(\tilde{u}, \delta^{(N)}), \tilde{u}, \delta^{(N)})) \leq 0, \quad j = 1, \dots, J. \end{aligned}$$

The Q -design problem (OC- Q_θ) can be approximated similarly. The decision variables in these problems (\tilde{u} and θ , respectively) can be regularized to avoid overfitting the training data $\delta^{(1)}, \dots, \delta^{(N)}$. The computational complexity of the sampled problem generally scales linearly with N , so relatively large sample sizes can often be used. Different sample sizes N_0, \dots, N_J can be used to approximate the different risk measures $\mathcal{R}_0, \dots, \mathcal{R}_J$.

VI. NUMERICAL EXAMPLE

In this section, we demonstrate the new capabilities developed in this paper through a numerical example. The example includes nonlinear dynamics, non-additive uncertainty and the CVaR $_\beta$ risk measure.

We consider a single-input, single-output system with dynamics

$$x_{t+1} = (1 - w(\delta_{t+1})/20)(x_t)_+^2 - u_t + w(\delta_{t+1}). \quad (7)$$

We assume the controller has perfect state information ($y_t = x_t$), so that we can directly compare control methods in isolation from the state estimation problem. The unforced dynamics have a stable equilibrium at the origin and an unstable equilibrium at unity. Our primary control objective is to maintain the state within the basin of attraction of the stable equilibrium, *i.e.*, to satisfy the constraint $c_1(x) \leq 0$, where

$$c_1(x) = \max\{x_1, \dots, x_T\} - 1.$$

We accept that this constraint may occasionally be violated due to the uncertain initial state and disturbances, but tolerate violations only if they are both infrequent and small. For this reason, we measure the risk of constraint violations by the conditional value-at-risk ($\mathcal{R}_1 = \text{CVaR}_\beta$) at confidence level $\beta = 0.95$. This risk measure addresses both the probability and (conditionally) expected magnitude of constraint violations.

As discussed in §II-B, $\text{CVaR}_\beta c_1(x)$ can be constrained by minimizing a scalar α subject to

$$\mathbf{E}(c_1(x) - \alpha)_+ \leq -\alpha(1 - \beta).$$

We would also like to minimize the total cost of control effort,

$$c_0(u, \delta) = p(\delta_1)|u_0| + \dots + p(\delta_T)|u_{T-1}|,$$

where $p(\delta_1), \dots, p(\delta_T) \geq 0$ are uncertain prices. We choose the cost risk measure $\mathcal{R}_0 = \mathbf{E}$.

The dynamics (7), while nonlinear, are nondecreasing in x_t and convex in (x_t, u_t) for all δ_{t+1} . (We assume $w(\cdot) \leq 20$, so the coefficient multiplying $(x_t)_+^2$ is nonnegative.) The cost function c_0 is convex in u for all δ , and the constraint function c_1 is convex and nondecreasing in each of x_1, \dots, x_T . By Theorem 1, therefore, the system is convex. It is also purifiable, since the controller has perfect state information and the dynamics are invertible in the disturbance $w(\delta_{t+1})$. The purified output at time t is $\xi_t(\delta_{0:t}) = w(\delta_t)$, with purifier

$$p_t(y_{0:t}, u_{0:t-1}) = \frac{y_t - (y_{t-1})_+^2 + u_t}{1 + (y_{t-1})_+^2/20}.$$

The inverse purifier q_t can be straightforwardly constructed from the system dynamics. Theorem 2 therefore applies, establishing a bijection between disturbance feedback policies and state feedback policies for this system. This gives our (infinite-dimensional, convex) control design problem:

$$\begin{aligned} & \underset{Q \in \mathcal{Q}, \alpha}{\text{minimize}} && \alpha + \mathbf{E} c_0(Q(\xi(\delta)), \delta) \\ & \text{subject to} && \mathbf{E}(c_1(\phi(Q(\xi(\delta)), \delta)) - \alpha)_+ \leq -\alpha(1 - \beta). \end{aligned} \quad (8)$$

We compare five suboptimal controllers for this system. The first is the optimal affine disturbance feedback policy, computed via the Q -design procedure. The second is a nonlinear Q -designed policy with piecewise quadratic basis policies tuned through trial and error. The third is the certainty-equivalent variant of open-loop MPC discussed in §V-A1; subproblems are solved under a single prediction of the exogenous input trajectory (in this case, its conditional expectation). The fourth is open-loop MPC with the same risk measures as Problem (8). The fifth is closed-loop MPC with the same risk measures and an affine disturbance feedback recourse model. We refer to the fourth and fifth controllers as open-loop and closed-loop *scenario* MPC, respectively, as is common in the literature. [51], [89]

Implementing each of these controllers involves numerically solving subproblems generated by Problem (8). As discussed in §V-B, we replace the expected values in these subproblems by sample averages over a training set of 1,000 sample exogenous input trajectories. We use a validation set of another 1,000 samples to tune parameters such as the basis policies in nonlinear Q -design and the MPC prediction horizon. We then compare the policies' performance in a test set of another 2,000 samples.

In simulations, the initial state has a half-normal distribution on $[1, \infty)$ with mean two. The disturbances $w(\delta_1), \dots, w(\delta_T)$ are identically Beta(2, 4) distributed, shifted and scaled to have mean zero and support $[-1, 2]$. Each price $p(\delta_t)$ of control effort is exponentially distributed with mean $2 + \cos(2\pi t/T)$;

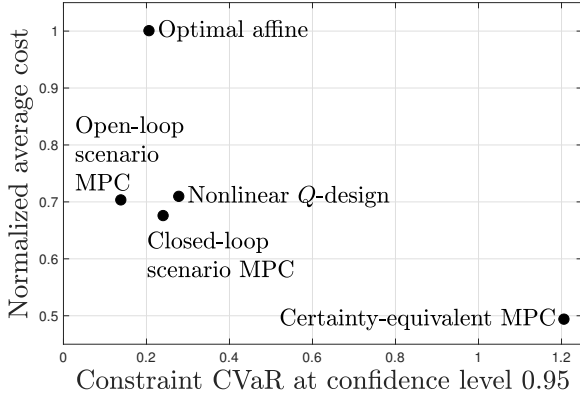


Fig. 2. Trade-offs between the cost of control effort and the risk of constraint violations in the test set of 2,000 sample exogenous input trajectories.

prices are lowest (on average) in the middle of the control horizon. The prices, disturbances and initial state are mutually independent.

Optimization is done in MATLAB using the Gurobi solver and the CVX modeling toolbox [90], which makes specifying our problems very easy. For example, the following code computes the optimal open-loop control input trajectory over N sample exogenous input trajectories stored in x_0 ($1 \times N$), w ($T \times N$) and p ($T \times N$):

```
variables a u(T,1)
expression x(T,N)
x = stateTrajectories(x0,u,w,f,T);
minimize( a + sum(c0(u,p))/N );
subject to
sum(pos(c1(x) - a))/N <= -a*(1 - b)
```

We wrote the vectorized function `stateTrajectories` to play the role of the input-state mapping ϕ , recursively building the state trajectory in each of the N Monte Carlo simulations from the dynamics function f . It typically speeds up CVX modeling by an order of magnitude compared to a loop over Monte Carlo runs. This function, along with all other code used in this paper, is online at [91].

Figure 2 shows the sample-average cost and constraint conditional value-at-risk of each controller in the test set of exogenous input trajectories. Table III shows the frequency with which each controller satisfies the constraint $x_t \leq 1$ in the test set. The five controllers strike different trade-offs between cost and robust constraint satisfaction. Certainty-equivalent MPC is aggressive, achieving low cost but poor robustness. The optimal affine controller satisfies constraints more robustly at the cost of increased control effort. The Q -designed nonlinear controller performs significantly better than the optimal affine controller, achieving similar cost and robustness to both of the scenario MPC variants. Including an affine recourse model in scenario MPC reduces cost slightly over open-loop scenario MPC, but also increases the risk of constraint violation. In this example, the robustness improvement of scenario MPC over certainty-equivalent MPC comes mainly from the use of samples, rather than the recourse model.

TABLE III
SAMPLE FREQUENCY OF $x_t \leq 1$

Controller	Frequency
Optimal affine	99.5%
Open-loop scenario MPC	99.3%
Nonlinear Q -design	98.9%
Closed-loop scenario MPC	98.7%
Certainty-equivalent MPC	74.6%

This numerical example demonstrates the key message of this paper: Suboptimal control methods that perform well for linear convex systems can be applied directly to nonlinear convex systems using the same software, and can perform similarly well. In this example, the Q -designed and predictive controllers show similar performance trends to the trends demonstrated for linear systems in [41], [42] and elsewhere. In particular, (a) scenario MPC is more robust than certainty-equivalent MPC, while achieving lower cost than affine controllers, and (b) nonlinear Q -designed controllers with good basis policies can be competitive with scenario MPC.

VII. CONCLUSION

In this paper, we explored the class of nonlinear systems for which optimal control design under uncertainty can be cast as a convex optimization problem. We adopted a flexible approach to the risks associated with constraint violations and excessive costs, accommodating both the robust and stochastic views. We showed that open-loop optimal control is convex for nonlinear systems with convex or concave dynamics, provided the dynamics, cost and constraint functions satisfy certain monotonicity properties. We then showed that under the same conditions, closed-loop control design can be reformulated as a convex program if, in addition, the measured outputs can be reversibly ‘purified’ of the influence of the past control inputs.

The practical value of these results is the guarantee that, for a class of nonlinear systems, the subproblems solved in various suboptimal control methods are convex. This removes concerns about solvers getting stuck in local minima, and enables the use of convenient modeling software and reliable, efficient solvers. We illustrated this numerically for two methods, the Q -design procedure and model predictive control.

There are a number of opportunities to extend this work. First, it would be interesting to survey applications in which nonlinear convex systems arise. We are aware of several examples, including cancer and HIV treatment scheduling [75], voltage control in electricity distribution networks [75], freeway traffic congestion management [76] and energy storage control [84], [85], but this list is likely incomplete.

Second, the requirement of purifiability appears to significantly restrict the uncertainty structures and nonlinearities admissible for Q -design. Exploring purifiability further could enable convex control design for a richer class of nonlinear systems. We note, however, that other methods, notably including model predictive control, do not require purifiability.

Third, the class of nonlinear systems for which optimal control is convex could be further explored. We characterized

only a subset of this class of systems; our method relied on the application of simple composition rules. While this method is constructive and compatible with convenient modeling software such as CVX [90] and YALMIP [92], it is not exhaustive.

APPENDIX A PROOF OF THEOREM 1

Our task is to show that the function

$$u \mapsto c_j((C(\delta)u + v(\delta), \eta(u, \delta)), u, \delta)$$

is convex for all δ . Here we have expressed the state trajectory as

$$x = \begin{bmatrix} x^{\text{aff}} \\ x^{\text{cvx}} \end{bmatrix} = \begin{bmatrix} C(\delta)u + v(\delta) \\ \eta(u, \delta) \end{bmatrix}.$$

We will define the mappings C , v and η shortly.

We recall that convexity is preserved under composition with an affine mapping and under the composition of a convex nondecreasing function with a convex function [83], and that c_j is nondecreasing in each element of x^{cvx} by assumption. It therefore suffices to show that each row of $\eta(\cdot, \delta)$ is convex for all δ .

To do this, we need the recursive definitions of C , v and η . They are initialized by

$$\begin{aligned} C_1(\delta_1) &= B_1(\delta_1) \\ v_1(\delta_{0:1}) &= A_1(\delta_1)w_0(\delta_0) + w_1(\delta_1) \\ \eta_1(u_0, \delta_{0:1}) &= h_1(w_0(\delta_0), h_0(\delta_0), u_0, \delta_1). \end{aligned}$$

For $t = 2, \dots, T$,

$$\begin{aligned} C_t(\delta_{1:t}) &= [A_t(\delta_t)C_{t-1}(\delta_{1:t-1}), B_t(\delta_t)] \\ v_t(\delta_{0:t}) &= A_t(\delta_t)v_{t-1}(\delta_{0:t-1}) + w_t(\delta_t) \\ \eta_t(u_{0:t-1}, \delta_{0:t}) &= h_t(C_{t-1}(\delta_{t-1})u_{0:t-2} + v_{t-1}(\delta_{0:t-1}), \\ &\quad \eta_{t-1}(u_{0:t-2}, \delta_{0:t-1}), u_{t-1}, \delta_t). \end{aligned}$$

The proof now proceeds by induction. At $t = 1$, η_1 is convex in u_0 by the convexity of h_1 . For the inductive step, we suppose that each row of η_{t-1} is convex in $u_{0:t-2}$. By assumption, each row of h_t is nondecreasing in x_{t-1}^{cvx} and convex. Therefore, η_t involves compositions of convex functions with the affine mapping $u_{0:t-2} \mapsto C_{t-1}(\delta_{t-1})u_{0:t-2} + v_{t-1}(\delta_{0:t-1})$ and of convex nondecreasing functions with convex functions, namely the rows of $u_{0:t-2} \mapsto \eta_{t-1}(u_{0:t-2}, \delta_{0:t-1})$. Applying our two composition rules, we see that each row of η_t is convex. This concludes the inductive step and the proof.

APPENDIX B PROOF OF THEOREM 2

Our task is to show that purifiability implies a bijection exists between causal π in y and causal Q in some purified output that depends on δ alone. To do this, we suppose the system is purifiable. We build a causal π in y from a given causal Q in the purified output

$$e = \begin{bmatrix} e_0 \\ e_1 \\ \vdots \\ e_{T-1} \end{bmatrix} = \begin{bmatrix} \xi_0(\delta_0) \\ \xi_1(\delta_{0:1}) \\ \vdots \\ \xi_{T-1}(\delta_{0:T-1}) \end{bmatrix} = \xi(\delta).$$

We recall that $e_0 = y_0$, always. The construction of π from Q is recursive:

$$\begin{aligned} \pi_0(y_0) &= Q_0(y_0) \\ \pi_t(y_{0:t}) &= Q_t(p_{0:t}(y_{0:t}, \pi_{0:t-1}(y_{0:t-1}))), \\ &\quad t = 1, \dots, T-1. \end{aligned} \tag{9}$$

Supposing this holds for $t = 0, \dots, \tau - 1$, at $t = \tau$ we have

$$\begin{aligned} u_\tau &= Q_\tau(e_{0:\tau}) \\ &= Q_\tau(p_{0:\tau}(y_{0:\tau}, u_{0:\tau-1})) \\ &= Q_\tau(p_{0:\tau}(y_{0:\tau}, \pi_{0:\tau-1}(y_{0:\tau-1}))) \\ &=: \pi_\tau(y_{0:\tau}). \end{aligned}$$

By induction, therefore, the recursion (9) holds for all $t = 0, \dots, T-1$. Given a causal π in y , a causal Q in e can be constructed similarly from the recursion

$$\begin{aligned} Q_0(e_0) &= \pi_0(e_0) \\ Q_t(e_{0:t}) &= \pi_t(q_{0:t}(e_{0:t}, Q_{0:t-1}(e_{0:t-1}))), \\ &\quad t = 1, \dots, T-1. \end{aligned}$$

This establishes the bijection and concludes the proof.

ACKNOWLEDGMENTS

The authors thank Anders Rantzer, Andy Ruina and Eilyan Bitar for their helpful comments, and gratefully acknowledge support from the National Science Foundation under grant 1711546.

REFERENCES

- [1] P. Whittle, *Optimization Over Time: Dynamic Programming and Stochastic Control*. Wiley, 1982.
- [2] P. Kumar and P. Varaiya, *Stochastic systems: Estimation, identification, and adaptive control*. SIAM, 1986.
- [3] A. Prekopa, *Stochastic Programming*. Kluwer Academic Publishers, 1995.
- [4] D. Bertsekas, *Dynamic Programming and Optimal Control*. Athena Scientific, 2005.
- [5] K. Astrom, *Introduction to Stochastic Control Theory*. Dover, 2006.
- [6] A. Shapiro, D. Dentcheva, and A. Ruszczyński, *Lectures on Stochastic Programming: Modeling and Theory*, ser. MPS-SIAM Series on Optimization. Philadelphia, PA: MPS-SIAM, 2009.
- [7] K. Zhou, J. Doyle, and K. Glover, *Robust and Optimal Control*. Prentice Hall, 1996.
- [8] A. Ben-Tal and A. Nemirovski, "Robust convex optimization," *Mathematics of Operations Research*, vol. 23, no. 4, pp. 769–805, 1998.
- [9] —, "Selected topics in robust convex optimization," *Mathematical Programming*, vol. 112, no. 1, pp. 125–158, 2008.
- [10] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski, *Robust Optimization*. Princeton University Press, 2009.
- [11] D. Bertsimas, D. Brown, and C. Caramanis, "Theory and applications of robust optimization," *SIAM Review*, vol. 53, no. 3, pp. 464–501, 2011.
- [12] P. Artzner, F. Delbaen, J. M. Eber, and D. Heath, "Thinking coherently," *Risk*, pp. 68–71, 1997.
- [13] —, "Coherent measures of risk," *Mathematical Finance*, vol. 9, no. 3, pp. 203–288, 1999.
- [14] R. T. Rockafellar and S. Uryasev, "Optimization of conditional value-at-risk," *Journal of risk*, vol. 2, pp. 21–42., 2000.
- [15] —, "Conditional value-at-risk for general loss distributions," *Journal of Banking and Finance*, vol. 26, pp. 1443–1471, 2002.
- [16] R. T. Rockafellar, "Coherent approaches to risk in optimization under uncertainty," *OR Tools and Applications: Glimpses of Future Technologies.*, pp. 38–61, 2007.
- [17] —, "Risk and utility in the duality framework of convex analysis," 2018.
- [18] R. Bellman, *Dynamic Programming*. Courier Dover Publications, 1957.

- [19] G. Newton, L. Gould, and J. Kaiser, *Analytical design of linear feedback controls*. Wiley, 1957.
- [20] R. Kalman, "Contributions to the theory of optimal control," *Boletin de la Sociedad Matematica Mexicana*, vol. 5, pp. 102–119, 1960.
- [21] D. Jacobson, "Optimal stochastic linear systems with exponential performance criteria and their relation to deterministic differential games," *IEEE Transactions on Automatic Control*, vol. 18, no. 2, pp. 124–131, 1973.
- [22] P. Whittle, "Risk-sensitive linear/quadratic/gaussian control," *Advances in Applied Probability*, vol. 13, no. 4, pp. 764–777, 1981.
- [23] D. Bertsekas and J. Tsitsiklis, "Neuro-dynamic programming: an overview," in *Conference on Decision and Control*, vol. 1, 1995, pp. 560–564.
- [24] R. Sutton and A. Barto, *Reinforcement learning: An introduction*. MIT Press, 1998.
- [25] T. Lillicrap, J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, "Continuous control with deep reinforcement learning," 2015, arXiv preprint 1509.02971.
- [26] V. Mnih, A. Badia, M. Mirza, A. Graves, T. Lillicrap, T. Harley, D. Silver, and K. Kavukcuoglu, "Asynchronous methods for deep reinforcement learning," in *International Conference on Machine Learning*, 2016.
- [27] D. De Fariis and B. Van Roy, "The linear programming approach to approximate dynamic programming," *Operations Research*, 2003.
- [28] W. Powell, *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. Wiley, 2007.
- [29] F. Lewis and D. Liu, *Reinforcement learning and approximate dynamic programming for feedback control*. Wiley, 2013.
- [30] C. Swamy and D. Shmoys, "Sampling-based approximation algorithms for multi-stage stochastic optimization," in *Foundations of Computer Science, 46th Annual IEEE Symposium on*, 2005, pp. 357–366.
- [31] A. Gupta, M. Pal, R. Ravi, and A. Sinha, *Approximation, Randomization and Combinatorial Optimization: Algorithms and Techniques*. Springer, 2005, ch. What about Wednesday? Approximation algorithms for multistage stochastic optimization.
- [32] C. Donohue and J. Birge, "The abridged nested decomposition method for multistage stochastic linear programs with relatively complete recourse," *Algorithmic Operations Research*, vol. 1, no. 1, 2006.
- [33] D. Kuhn, W. Wiesemann, and A. Georghiou, "Primal and dual linear decision rules in stochastic and robust optimization," *Mathematical Programming*, vol. 130, no. 1, pp. 177–209, 2011.
- [34] A. Ben-Tal, S. Boyd, and A. Nemirovski, "Extending scope of robust optimization: Comprehensive robust counterparts of uncertain problems," *Mathematical Programming*, vol. 107, no. 1-2, pp. 63–89, 2006.
- [35] H. Beyer and B. Sendhoff, "Robust optimization—a comprehensive survey," *Computer methods in applied mechanics and engineering*, vol. 196, no. 33-34, pp. 3190–3218, 2007.
- [36] X. Chen and Y. Zhang, "Uncertain linear programs: Extended affinely adjustable robust counterparts," *Operations Research*, vol. 57, no. 6, pp. 1469–1482, 2009.
- [37] V. Kucera, *Discrete linear control: the polynomial equation approach*. John Wiley & Sons, Inc., 1980.
- [38] C. Desoer and C. Lin, "Non-linear unity-feedback systems and q-parametrization," *International Journal of Control*, vol. 40, no. 1, pp. 37–51, 1984.
- [39] V. Anantharam and C. Desoer, "On the stabilization of nonlinear systems," *IEEE Transactions on Automatic Control*, vol. 29, no. 6, pp. 569–572, 1984.
- [40] S. Boyd and C. Barratt, *Linear controller design: Limits of performance*. Stanford University, 1991.
- [41] J. Skaf and S. Boyd, "Nonlinear q-design for convex stochastic control," *IEEE Transactions on Automatic Control*, vol. 54, no. 10, pp. 2426–2430, 2009.
- [42] —, "Design of affine controllers via convex optimization," *IEEE Transactions on Automatic Control*, vol. 55, no. 11, pp. 2476–2487, 2010.
- [43] D. Mayne, J. Rawlings, C. Rao, and P. Scokaert, "Constrained model predictive control: Stability and optimality," *Automatica*, vol. 36, pp. 789–814, 2000.
- [44] A. Bemporad, F. Borrelli, and M. Morari, "Model predictive control based on linear programming: the explicit solution," *IEEE Transactions on Automatic Control*, vol. 47, no. 12, pp. 1974–1985, 2002.
- [45] Y. Wang and S. Boyd, "Fast model predictive control using online optimization," *IEEE Transactions on Control Systems Technology*, vol. 18, no. 2, pp. 267–278, 2010.
- [46] A. Schwarm and M. Nikolaou, "Chance-constrained model predictive control," *AIChE Journal*, vol. 45, no. 8, pp. 1743–1752, 1999.
- [47] F. Oldewurtel, C. Jones, and M. Morari, "A tractable approximation of chance constrained stochastic mpc based on affine disturbance feedback," in *Conference on Decision and Control*, 2008.
- [48] L. Blackmore, M. Ono, A. Bektassov, and B. Williams, "A probabilistic particle-control approximation of chance-constrained stochastic predictive control," *Robotics, IEEE Transactions on*, vol. 26, no. 3, pp. 502–517, 2010.
- [49] M. Cannon, B. Kouvaritakis, S. Rakovic, and Q. Cheng, "Stochastic tubes in model predictive control with probabilistic constraints," *IEEE Transactions on Automatic Control*, vol. 56, no. 1, pp. 194–200, 2011.
- [50] E. Cinquemani, M. Agarwal, D. Chatterjee, and J. Lygeros, "Convexity and convex approximations of discrete-time stochastic control problems with constraints," *Automatica*, vol. 47, no. 9, pp. 2082–2087, 2011.
- [51] G. Schildbach, L. Fagiano, C. Frei, and M. Morari, "The scenario approach for stochastic model predictive control with bounds on closed-loop constraint violations," *Automatica*, vol. 50, no. 12, pp. 3009–3018, 2014.
- [52] K. J. Kircher and K. M. Zhang, "Sample-average model predictive control of uncertain linear systems," in *Conference on Decision and Control*, 2016, pp. 6234–6239.
- [53] P. Campo and M. Morari, "Robust model predictive control," in *American Control Conference (ACC)*, 1987, pp. 1021–1026.
- [54] L. Magni, G. De Nicolao, R. Scattolini, and F. Allgower, "Robust model predictive control for nonlinear discrete-time systems," *International Journal of Robust and Nonlinear Control*, vol. 13, no. 3, pp. 229–246, 2003.
- [55] D. Mayne, M. Seron, and S. Rakovic, "Robust model predictive control of constrained linear systems with bounded disturbances," *Automatica*, vol. 41, no. 2, pp. 219–224, 2005.
- [56] P. Goulart and E. Kerrigan, "Output feedback receding horizon control of constrained systems," *International Journal of Control*, vol. 80, no. 1, pp. 8–20, 2007.
- [57] G. Calafiore and L. Fagiano, "Robust model predictive control via scenario optimization," *IEEE Transactions on Automatic Control*, vol. 58, no. 1, pp. 219–224, 2013.
- [58] R. Kalman, "A new approach to linear filtering and prediction problems," *Journal of Basic Engineering*, vol. 82, no. 1, pp. 35–45, 1960.
- [59] R. Kalman and R. Bucy, "New results in linear filtering and prediction theory," *Journal of basic engineering*, vol. 83, no. 1, pp. 95–108, 1961.
- [60] R. Kopp and R. Orford, "Linear regression applied to system identification for adaptive control systems," *AIAA Journal*, vol. 1, no. 10, pp. 2300–2306, 1963.
- [61] H. Cox, "On the estimation of state variables and parameters for noisy dynamic systems," *IEEE Transactions on Automatic Control*, vol. 9, no. 1, pp. 5–12, 1964.
- [62] L. Ljung, "Asymptotic behavior of the extended kalman filter as a parameter estimator for linear systems," *IEEE Transactions on Automatic Control*, vol. 24, no. 1, pp. 36–50, 1979.
- [63] E. Wan and R. Van Der Merwe, "The unscented Kalman filter for nonlinear estimation," in *The IEEE Adaptive Systems for Signal Processing, Communications, and Control Symposium*, 2000, pp. 153–158.
- [64] R. Van Der Merwe and E. Wan, "The square-root unscented kalman filter for state and parameter-estimation," in *Acoustics, Speech, and Signal Processing, IEEE International Conference on*, vol. 6, 2001, pp. 3461–3464.
- [65] G. Kitagawa, "Monte carlo filter and smoother for non-gaussian nonlinear state space models," *Journal of computational and graphical statistics*, vol. 5, no. 1, pp. 1–25, 1996.
- [66] M. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, "A tutorial on particle filters for online nonlinear/non-gaussian bayesian tracking," *IEEE Transactions on signal processing*, vol. 50, no. 2, pp. 174–188, 2002.
- [67] D. Youla, J. Bongiorno, and H. Jabr, "Modern wiener-hopf design of optimal controllers part i: The single-input-output case," *IEEE Transactions on Automatic Control*, vol. 21, no. 1, pp. 3–13, 1976.
- [68] D. Youla, H. Jabr, and J. Bongiorno, "Modern wiener-hopf design of optimal controllers—part ii: The multivariable case," *IEEE Transactions on Automatic Control*, vol. 21, no. 3, pp. 319–338, 1976.
- [69] J. Wu and S. Lall, "Nonlinear youla parametrization and information constraints for decentralized control," in *American Control Conference (ACC)*, 2010, pp. 5614–5619.
- [70] D. Van Hessem and O. Bosgra, "A conic reformulation of model predictive control including bounded and stochastic disturbances under state and input constraints," in *Conference on Decision and Control*, 2002.
- [71] J. Lofberg, "Approximations of closed-loop minimax mpc," in *Conference on Decision and Control*, 2003.

- [72] A. Ben-Tal, A. Goryashko, E. Guslitzer, and A. Nemirovski, "Adjustable robust solutions of uncertain linear programs," *Mathematical Programming*, vol. 99, no. 2, pp. 351–376, 2004.
- [73] P. Goulart, E. Kerrigan, and J. Maciejowski, "Optimization over state feedback policies for robust control with constraints," *Automatica*, 2006.
- [74] C. Wang, C. Ong, and M. Sim, "Model predictive control using segregated disturbance feedback," *IEEE Transactions on Automatic Control*, vol. 55, no. 4, pp. 831–840, 2010.
- [75] A. Rantzer and B. Bernhardsson, "Control of convex-monotone systems," in *Conference on Decision and Control*, 2014, pp. 2378–2383.
- [76] M. Schmitt, C. Ramesh, P. Goulart, and J. Lygeros, "Convex, monotone systems are optimally operated at steady-state," in *American Control Conference (ACC)*, 2017, pp. 2662–2667.
- [77] G. Calafiore and M. Campi, "Uncertain convex programs: Randomized solutions and confidence levels," *Mathematical Programming*, vol. 102, no. 1, pp. 25–46, 2005.
- [78] M. Campi and S. Garatti, "The exact feasibility of randomized solutions of uncertain convex programs," *SIAM Journal on Optimization*, vol. 19, no. 3, pp. 1211–1230, 2008.
- [79] G. Calafiore, "Random convex programs," *SIAM Journal on Optimization*, vol. 20, no. 6, pp. 3427–3464, 2010.
- [80] G. Schildbach, L. Fagiano, and M. Morari, "Randomized solutions to convex programs with multiple chance constraints," *SIAM Journal on Optimization*, vol. 23, no. 4, pp. 2479–2051, 2013.
- [81] X. Zhang, S. Grammatico, G. Schildbach, P. Goulart, and J. Lygeros, "On the sample size of random convex programs with structured dependence on the uncertainty," *Automatica*, vol. 60, pp. 182–188, 2015.
- [82] M. Campi and S. Garatti, "Wait-and-judge scenario optimization," *Mathematical Programming*, vol. 167, no. 1, pp. 155–189, 2018.
- [83] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.
- [84] R. Khodabakhsh and S. Sirouspour, "Optimal control of energy storage in a microgrid by minimizing conditional value-at-risk," *IEEE Transactions on Sustainable Energy*, vol. 7, no. 3, pp. 1264–1273, 2016.
- [85] Y. Xu and L. Tong, "Optimal operation and economic value of energy storage at consumer locations," *IEEE Transactions on Automatic Control*, vol. 62, no. 2, pp. 792–807, 2017.
- [86] A. Kleywegt, A. Shapiro, and T. Homem-de Mello, "The sample average approximation method for stochastic discrete optimization," *SIAM Journal on Optimization*, vol. 12, no. 2, pp. 479–502, 2001.
- [87] T. Homem-de Mello and G. Bayraksan, "Monte carlo sampling-based methods for stochastic optimization," *Surveys in Operations Research and Management Science*, vol. 19, no. 1, pp. 56–85, 2014.
- [88] S. Kim, R. Pasupathy, and S. G. Henderson, "A guide to sample average approximation," in *Handbook of Simulation Optimization*. Springer New York, 2015, pp. 207–243.
- [89] G. Calafiore and L. Fagiano, "Stochastic model predictive control of LPV systems via scenario optimization," *Automatica*, vol. 49, no. 6, pp. 1861–1866, 2013.
- [90] M. Grant and S. Boyd, *Graph Implementations for Nonsmooth Convex Programs, Recent Advances in Learning and Control (a tribute to M. Vidyasagar)*, ser. Lecture Notes in Control and Information Sciences, V. Blondel, S. Boyd, and H. Kimura, Eds. Springer, 2008.
- [91] [Online]. Available: <https://github.com/kevinjkircher>
- [92] J. Lofberg, "Yalmip: A toolbox for modeling and optimization in matlab," in *Computer Aided Control Systems Design, IEEE International Symposium on*, 2004, pp. 284–289.



K. Max Zhang is an Associate Professor in the Sibley School of Mechanical and Aerospace Engineering at Cornell University, where he directs the Energy and Environment Research Laboratory. He received the Ph.D. degree in mechanical engineering from the University of California, Davis, in 2004. He was a visiting scientist at the United States Environmental Protection Agency (EPA) in 2000, 2002 and 2010–2012. Dr. Zhang received the EPA's Scientific and Technological Achievement Award and Cornell University's Engaged Scholar Prize.



Kevin J. Kircher (M'15) is a mechanical engineering Ph.D. student at Cornell University. He received the B.S. degree in applied mathematics and physics with high honors from the University of Wisconsin-Milwaukee in 2008 and the M.Eng. degree in engineering physics from Cornell University in 2009. He then worked as a research associate at Lawrence Berkeley National Laboratory. His research interests include optimization and control under uncertainty, with applications to heating, cooling and power systems.