

Constrained Stochastic LQC: A Tractable Approach

Dimitris Bertsimas and David B. Brown

Abstract—Despite the celebrated success of dynamic programming for optimizing quadratic cost functions over linear systems, such an approach is limited by its inability to tractably deal with even simple constraints. In this paper, we present an alternative approach based on results from robust optimization to solve the stochastic linear-quadratic control (SLQC) problem. In the unconstrained case, the problem may be formulated as a semidefinite optimization problem (SDP). We show that we can reduce this SDP to optimization of a convex function over a scalar variable followed by matrix multiplication in the current state, thus yielding an approach that is amenable to closed-loop control and analogous to the Riccati equation in our framework. We also consider a tight, second-order cone (SOCP) approximation to the SDP that can be solved much more efficiently when the problem has additional constraints. Both the SDP and SOCP are tractable in the presence of control and state space constraints; moreover, compared to the Riccati approach, they provide much greater control over the stochastic behavior of the cost function when the noise in the system is distributed normally.

Index Terms—Control with constraints, linear-quadratic control, robust optimization, semidefinite optimization.

I. INTRODUCTION

THE theory of dynamic programming, while conceptually elegant, is computationally impractical for all but a few special cases of system dynamics and cost functions. One of the notable triumphs of dynamic programming is its success with stochastic linear systems and quadratic cost functions (stochastic linear-quadratic control—SLQC). It is easily shown (e.g., [4]) in this case that the cost-to-go functions are quadratic in the state, and therefore the resulting optimal controls are linear in the current state. As a result, solving Bellman's equation in this case is tantamount to finding appropriate gain matrices, and these gain matrices are described by the well-known Riccati equation [19].

This success, however, has some limitations. In particular, Bellman's equation in the SLQC has tractability issues with even the simplest of constraints on either the control or state vectors. It is not difficult to find applications that demand constraints on the controls or state. Bertsimas and Lo [5] describe the dynamics of an optimal share-purchasing policy for stockholders. The unconstrained policy based on the Riccati equation requires the investor to purchase and *sell* shares, which is clearly absurd. This can be mitigated by a nonnegativity constraint on the control, which causes the cost-to-go function to be-

come piecewise quadratic with an exponential number of pieces. Thus, a very simple constraint destroys the tractability of this approach. Much of the current literature (e.g., [17] and [18]) derives necessary conditions for optimality for simple control constraints but does not explicitly describe solution methods.

A further drawback of the Riccati approach from dynamic programming is that it only deals with the expected value of the resulting cost. In many cases, we may wish to know more information about the distribution of the cost function (e.g., cases in which we want to provide a probabilistic level of protection guaranteeing some system performance).

In this paper, we propose an alternative approach to the SLQC problem. Rather than attempting to solve Bellman's equation, we exploit relatively new results from robust optimization to propose an alternative solution technique for SLQC. Our approach has the following advantages over the traditional dynamic programming approach.

- 1) It can tractably handle a variety of constraints on both the control and state vectors.
- 2) It admits a probabilistic description of the resulting cost, allowing us to understand and control the system cost distribution.
- 3) In the unconstrained case, its complexity is not much more than the complexity of linear feedback (i.e., the Riccati approach). In particular, optimal policies in this case may be computed by optimizing a convex function over a scalar, then multiplying the initial state by appropriate matrices.

Our approach is based on techniques from robust optimization. Although the use of convex optimization techniques is common in the control literature (see, e.g., [7], [9], [13], and [14]), we believe our methodology is a new one. Chen and Zhou [8] provide an elegant solution to the SLQC problem with conic control constraints, but their solution is limited to a scalar-valued state variable and homogeneous system dynamics. Our approach here is more general. We emphasize that we are *not* proposing a solution for robust control (see, e.g., [21] for a start to the vast literature on the subject); rather, we are proposing an approach to the SLQC with the conceptual framework of robust optimization as a guide.

The structure of this paper is as follows. In Section II, we present a description of the SQLC problem, as well as the currently known results from dynamic programming and a conceptual description of our methodology. In addition, we provide background for the robust optimization results we will later use. In Section III, we develop our approach for the unconstrained SLQC problem. This approach is based on semidefinite programming (SDP) and robust quadratic programming results from Ben-Tal and Nemirovski [3]. We further show that this SDP has a very special structure that allows us to derive a closed-loop control law suitable for real-time applications. Unfortunately, in the presence of constraints, this simplification no

Manuscript received September 15, 2004; revised August 3, 2005 and August 14, 2006. Recommended by Associate Editor C. D. Charalambous.

D. Bertsimas is with the Sloan School of Management and Operations Research Center, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: dbertsim@mit.edu).

D. B. Brown is with the Fuqua School of Business, Duke University, Durham, NC 27705 USA (e-mail: dbbrown@duke.edu).

Digital Object Identifier 10.1109/TAC.2007.906182

longer applies, and the complexity of solving the SDP is impractical in a closed-loop setting. This motivates us to simplify the SDP, which we do in Section IV. Here we use recent results from robust conic optimization developed by Bertsimas and Sim [6] to develop a tight SOCP approximation that is far easier to solve. We then show in Section V how this approach admits various constraints and performance guarantees. These constraints may be deterministic constraints on the control or probabilistic guarantees on the state and objective function. In Section VI, we show that a particular model for imperfect state information fits into the framework already developed, and in Section VII, we provide computational results. Section VIII concludes this paper.

II. PROBLEM STATEMENT AND PRELIMINARIES

Throughout this paper, we will work with discrete-time stochastic linear systems of the form

$$\mathbf{x}_{k+1} = \mathbf{A}_k \mathbf{x}_k + \mathbf{B}_k \mathbf{u}_k + \mathbf{C}_k \mathbf{w}_k, \quad k = 0, \dots, N-1 \quad (1)$$

where $\mathbf{x}_k \in \mathbb{R}^{n_x}$ is a state vector, $\mathbf{u}_k \in \mathbb{R}^{n_u}$ is a control vector, and $\mathbf{w}_k \in \mathbb{R}^{n_w}$ is a disturbance vector (an unknown quantity). We assume throughout that the matrices $\mathbf{A}_k \in \mathbb{R}^{n_x \times n_x}$, $\mathbf{B}_k \in \mathbb{R}^{n_x \times n_u}$, and $\mathbf{C}_k \in \mathbb{R}^{n_x \times n_w}$ are known exactly.

It is desired to control the system in question in a way that keeps the cost function

$$J(\mathbf{x}_0, \mathbf{u}, \mathbf{w}) = \sum_{k=1}^N (\mathbf{x}_k^T \mathbf{Q}_k \mathbf{x}_k + 2\mathbf{q}_k^T \mathbf{x}_k) + \sum_{k=0}^{N-1} (\mathbf{u}_k^T \mathbf{R}_k \mathbf{u}_k + 2\mathbf{r}_k^T \mathbf{u}_k) \quad (2)$$

as small as possible. Here we will assume $\mathbf{Q}_k \succeq \mathbf{0}$, $\mathbf{R}_k \succ \mathbf{0}$, and, again, that the data \mathbf{Q}_k , \mathbf{q}_k , \mathbf{R}_k , and \mathbf{r}_k , are known exactly. We are also using the shorthand \mathbf{u} and \mathbf{w} to denote the entire vector of controls and disturbances, i.e.,

$$\mathbf{u}^T = [\mathbf{u}_0^T \quad \mathbf{u}_1^T \quad \dots \quad \mathbf{u}_{N-1}^T] \quad (3)$$

$$\mathbf{w}^T = [\mathbf{w}_0^T \quad \mathbf{w}_1^T \quad \dots \quad \mathbf{w}_{N-1}^T]. \quad (4)$$

Finally, our convention will be for the system to be in some initial state \mathbf{x}_0 . Unless otherwise stated, we assume this initial state is also known exactly.

Note that (2) is an uncertain quantity, as it depends on the realization of \mathbf{w} , which is unknown. Most approaches assume \mathbf{w} is a random variable possessing some distributional properties and proceed to minimize (2) in an expected value sense. We now survey the traditional approach to this problem.

A. The Traditional Approach: Bellman's Recursion

The dynamic programming approach requires a few distributional assumptions on the disturbance vectors. Typically, it is assumed that the \mathbf{w}_k are independent, and independent of both \mathbf{x}_k and \mathbf{u}_k . Moreover, we have $\mathbb{E}[\mathbf{w}_k] = \mathbf{0}$, and \mathbf{w}_k has finite second moment. For this derivation, we will assume $\mathbf{q}_k = \mathbf{0}$, $\mathbf{r}_k = \mathbf{0}$, and $\mathbf{C}_k = \mathbf{I}$ for ease of notation, but the result holds more generally after some simple manipulations. Modifications of some

of the distributional assumptions (such as nonzero mean, correlations) are also possible, but we do not detail them here.

The literature on this subject is vast, and the problem is well understood. The main result is that the expected cost-to-go functions $J_k(\mathbf{x}_k)$ defined by

$$\begin{aligned} J_N(\mathbf{x}_N) &= \mathbf{x}_N^T \mathbf{Q}_N \mathbf{x}_N \\ J_k(\mathbf{x}_k) &= \mathbf{x}_k^T \mathbf{Q}_k \mathbf{x}_k \\ &\quad + \min_{\mathbf{u}_k} \mathbb{E}[\mathbf{u}_k^T \mathbf{R}_k \mathbf{u}_k + J_{k+1}(\mathbf{A}_k \mathbf{x}_k + \mathbf{B}_k \mathbf{u}_k + \mathbf{w}_k)] \end{aligned} \quad (5)$$

are quadratic in the state \mathbf{x}_k . Thus, it follows that the optimal policy is *linear* in the current state. In particular, one can show (see, e.g., [4]) that the optimal control \mathbf{u}_k^* is given by

$$\mathbf{u}_k^* = \mathbf{L}_k \mathbf{x}_k$$

where $\mathbf{L}_k = -(\mathbf{B}_k^T \mathbf{K}_{k+1} \mathbf{B}_k + \mathbf{R}_k)^{-1} \mathbf{B}_k^T \mathbf{K}_{k+1} \mathbf{A}_k$ and \mathbf{K}_k are symmetric, positive semidefinite matrices computed recursively. The fact that the recursion given in (5) works so well (from a complexity standpoint) is quite particular to the case of linear systems and quadratic costs. For more arbitrary systems or cost functions such an approach is, in general, intractable.

A more troubling difficulty, however, is that even with the same system and cost function, this approach explodes computationally with ostensibly simple constraints, such as $\mathbf{u}_k \geq \mathbf{0}$. For instance, the cost-to-go function (5) in this case becomes piecewise quadratic with an exponential (in N) number of pieces.

Of course, one way to suboptimally handle this issue is to apply Lagrangian duality techniques to the constraints. For example, in the case of quadratic constraints on the control vectors, i.e., $\mathbf{u}_k^T \mathbf{P}_k \mathbf{u}_k \leq \rho_k^2$, one may relax the constraints and then maximize over a dual vector \mathbf{p} . In particular, the cost-to-go functions now have the form

$$\begin{aligned} J_N(\mathbf{x}_N) &= \mathbf{x}_N^T \mathbf{Q}_N \mathbf{x}_N \\ J_k(\mathbf{x}_k) &= \max_{\mathbf{p}_k \geq 0} f_k(\mathbf{p}_k, \mathbf{x}_k) \end{aligned} \quad (6)$$

where the dual functions $f_k(\mathbf{p}_k, \mathbf{x}_k)$ have the form

$$\begin{aligned} f_k(\mathbf{p}_k, \mathbf{x}_k) &= \mathbf{x}_k^T \mathbf{Q}_k \mathbf{x}_k + \min_{\mathbf{u}_k} \mathbb{E}[\mathbf{u}_k^T \mathbf{R}_k \mathbf{u}_k + \mathbf{p}_k (\rho_k^2 - \mathbf{u}_k^T \mathbf{P}_k \mathbf{u}_k) \\ &\quad + J_{k+1}(\mathbf{A}_k \mathbf{x}_k + \mathbf{B}_k \mathbf{u}_k + \mathbf{w}_k)]. \end{aligned}$$

From here, one approach to solving (6) suboptimally is to select a priori a dual vector \mathbf{p} and then apply the Riccati equation as usual. An optimal solution, however, relies on computation of the optimal dual vector \mathbf{p}^* , which, in general, is difficult and destroys the quadratic form of the cost-to-go functions.

Thus, the traditional, dynamic programming approach can be solved very rapidly with linear feedback in the unconstrained case but becomes, for large-scale problems, impossible to solve optimally when the constraints are included. This is a very unfavorable property of the DP approach, and it is in direct contrast to the field of convex optimization, whose problem instances are quite robust (in terms of complexity) to perturbations in the constraint structure. Our approach, which we now detail, will leverage this useful property of convex optimization.

B. A Tractable Approach: Overview

The traditional approach above is not amenable to problem changes such as the addition of constraints for two primary reasons.

- 1) *Complexity of distributional calculations.* Computing the expectation in (5), except for very special cases, is cumbersome computationally.
- 2) *Intractability of Bellman's recursion.* The recursion in (5) requires us, when computing the current control, to have advance knowledge of *all future controls for all possible future states*, even states that are extraordinarily improbable. While this recursion is an elegant idea conceptually, it is not well suited to computation because the number of possible future states grows so rapidly with problem size.

We propose the following approach, which circumvents these difficulties.

- a) Given our current state \mathbf{x}_0 and problem data, we consider the entire control and disturbance vectors $\mathbf{u} \in \mathbb{R}^{N \cdot n_u}$, $\mathbf{w} \in \mathbb{R}^{N \cdot n_w}$, respectively, as in (3) and (4).
- b) We do not assume a particular distribution for \mathbf{w} . Assume only that \mathbf{w} belongs within some “reasonable” uncertainty set. In particular, assume \mathbf{w} belongs to some norm-bounded set

$$\mathcal{W}_\gamma = \{\mathbf{w} \mid \|\mathbf{w}\|_2 \leq \gamma\} \quad (7)$$

parameterized by $\gamma \geq 0$.¹

- c) Discard the notion of Bellman's recursion. Instead, do the best we can for all possible disturbances within \mathcal{W}_γ . That is, rather than computing controls for *every* possible state realization, we simply choose a control vector for the remaining stages that performs best for the most pessimistic disturbance within this “reasonable” uncertainty set. Specifically, we search for an optimal control \mathbf{u}^* to the problem

$$\min_{\mathbf{u} \in \mathbb{R}^{N \cdot n_u}} \max_{\mathbf{w} \in \mathcal{W}_\gamma} J(\mathbf{x}_0, \mathbf{u}, \mathbf{w}). \quad (8)$$

Of course, this brings up the issue of open-loop versus closed-loop control. At first glance, this approach appears to be an open-loop method only. We can, however, compute a solution \mathbf{u}^* to (8), take the first n_u components, and apply this as the current control. After a new state observation, we can repeat the calculation in (8) with the updated problem data (most of this updating can be done offline). The only issue is that the routine for solving (8) be computationally simple enough for the application at hand. The complexity of these solution procedures will indeed be a central issue for much of the remaining discussion.

Note that the model in (8) is similar in spirit to the approach of H^∞ control (e.g., [2]) in that it is worst case over a deterministic uncertainty set. In contrast to H^∞ control, however, our methodology explicitly relies on new results in robust optimization. In particular, our approach has the following properties.

¹If we wish instead to have $\mathbf{w} \in \{\mathbf{w} \mid \mathbf{w}^T \Sigma^{-1} \mathbf{w} \leq \gamma^2\}$, where $\Sigma \succ \mathbf{0}$, then we may rescale coordinates and obtain a problem of the same form. Note that the statistical appropriateness of ellipsoids and their explicit construction is not the subject of this paper, but the interested reader may see Paganini [11] for uncertainty set modelling for the case of white noise.

- It is tractable, even in the presence of control and state-space constraints.
- We solve a *deterministic* problem (8) to compute an optimal solution \mathbf{u}_γ^* . Thus far, we have not discussed probability in any way. Nonetheless, our approach is amenable to examining how good \mathbf{u}_γ^* is when the disturbances \mathbf{w}_k , rather being chosen in an adversarial manner from an ellipsoid, instead obey a probability distribution. We show in Theorem 8 that under normality for \mathbf{w}_k , the solution \mathbf{u}_γ^* satisfies very strong probabilistic guarantees. In other words, when nature gives rise to disturbances that are bounded, we solve the problem optimally. When, on the other hand, nature gives rise to disturbances that do not satisfy $\|\mathbf{w}\| \leq \gamma$, we still show strong probabilistic guarantees on the performance of \mathbf{u}_γ^* .
- In the unconstrained case, it yields an efficient control law that is linear in the current state after a simple, scalar optimization procedure. In addition, for $\gamma = 0$, we recover the traditional (Riccati) solution, whereas, for $\gamma > 0$, we have a family of increasingly conservative approaches.

To solve (8), we will utilize a number of results from robust optimization, which we now describe.

C. Results From Robust Quadratic Optimization Over Ellipsoids

We will leverage some robust quadratic programming results popularized by Ben-Tal and Nemirovski [3]. In particular, they consider the conic quadratic constraint

$$\|\mathbf{A}\mathbf{x} + \mathbf{b}\|_2 \leq \mathbf{c}^T \mathbf{x} + d$$

when the data $(\mathbf{A}, \mathbf{b}, \mathbf{c}, d)$ are uncertain and known only to belong to some bounded uncertainty set \mathcal{U} . The goal of robust quadratic programming is to optimize over the set of all \mathbf{x} such that the constraint holds for all possible values of the data within the set \mathcal{U} . In other words, we desire to find \mathbf{x} such that

$$\|\mathbf{A}\mathbf{x} + \mathbf{b}\|_2 \leq \mathbf{c}^T \mathbf{x} + d \quad \forall (\mathbf{A}, \mathbf{b}, \mathbf{c}, d) \in \mathcal{U}.$$

Ben-Tal and Nemirovski show that in the case of an ellipsoidal uncertainty set, the problem of optimizing over an uncertain conic quadratic inequality may be solved tractably using semidefinite programming. This turns out also to be the case for (8). To this end, we will need the following two classical results, proofs of which may be found in [3], among others. First, we have the *Schur complement lemma*.

Lemma 1: Let

$$\mathbf{A} = \begin{bmatrix} \mathbf{B} & \mathbf{C}^T \\ \mathbf{C} & \mathbf{D} \end{bmatrix}$$

where $\mathbf{B} \succ \mathbf{0}$. Then \mathbf{A} is positive (semi) definite if and only if the matrix $\mathbf{D} - \mathbf{C}\mathbf{B}^{-1}\mathbf{C}^T$ is positive (semi) definite.

In addition, we have the *S-lemma*.

Lemma 2: Let \mathbf{A}, \mathbf{B} be symmetric $n \times n$ matrices and assume that the quadratic inequality

$$\mathbf{x}^T \mathbf{A} \mathbf{x} \geq 0$$

is strictly feasible. Then the minimum value of the problem

$$\begin{aligned} & \text{minimize} && \mathbf{x}^T \mathbf{B} \mathbf{x} \\ & \text{subject to} && \mathbf{x}^T \mathbf{A} \mathbf{x} \geq 0 \end{aligned}$$

is nonnegative if and only if there exists a $\lambda \geq 0$ such that $\mathbf{B} - \lambda \mathbf{A} \succeq \mathbf{0}$.

D. Results From Robust Conic Optimization Over Norm-Bounded Sets

To improve the complexity of solving (8) when we have constraints, we will utilize recent results from robust conic optimization results due to Bertsimas and Sim [6]. This approach is a relaxation of the exact min-max approach but is computationally less complex and leads to a unified probability bound across a variety of conic optimization problems. We survey the main ideas and developments here.

Bertsimas and Sim use the following model for data uncertainty:

$$\tilde{\mathbf{D}} = \mathbf{D}^0 + \sum_{j \in N} \Delta \mathbf{D}^j \tilde{z}_j$$

where \mathbf{D}^0 is the nominal data value and $\Delta \mathbf{D}^j$ are data perturbations. The \tilde{z}_j are random variables with mean zero and independent, identical distributions. The goal is to find a policy \mathbf{x} such that a given constraint is ‘‘robust feasible,’’ i.e.,

$$\max_{\tilde{\mathbf{D}} \in \mathcal{U}_\Omega} f(\mathbf{x}, \tilde{\mathbf{D}}) \leq 0 \quad (9)$$

where

$$\mathcal{U}_\Omega = \left\{ \mathbf{D}^0 + \sum_{j \in N} \Delta \mathbf{D}^j u_j \mid \|\mathbf{u}\| \leq \Omega \right\}. \quad (10)$$

For our purposes, we typically use the Euclidean norm on \mathbf{u} , as it is self-dual, but many other choices for the norm may be tractably used [6]. We operate under some restrictions on the function $f(\mathbf{x}, \mathbf{D})$.²

Assumption 1: The function $f(\mathbf{x}, \mathbf{D})$ satisfies the following.

- $f(\mathbf{x}, \mathbf{D})$ is convex in \mathbf{D} for all $\mathbf{x} \in \mathbb{R}^n$.
- $f(\mathbf{x}, k\mathbf{D}) = kf(\mathbf{x}, \mathbf{D})$ for all $k \geq 0, \mathbf{D}, \mathbf{x} \in \mathbb{R}^n$.

One of the central ideas of [6] is to linearize the model of robustness as follows:

$$\max_{(\mathbf{u}, \mathbf{v}) \in \mathcal{V}_\Omega} f(\mathbf{x}, \mathbf{D}^0) + \sum_{j \in N} \{f(\mathbf{x}, \Delta \mathbf{D}^j) u_j + f(\mathbf{x}, -\Delta \mathbf{D}^j) v_j\} \leq 0 \quad (11)$$

where

$$\mathcal{V}_\Omega = \left\{ (\mathbf{u}, \mathbf{v}) \in \mathbb{R}_+^{2|N|} \mid \|\mathbf{u} + \mathbf{v}\| \leq \Omega \right\}.$$

²In [6], the authors assume the function is concave in the data. For our purposes, convexity is more convenient. All results follow up to sign changes, and we report them accordingly.

In the framework developed thus far, (11) turns out to be a relaxation of (9), i.e., we have the following.

Proposition 1 (Bertsimas–Sim):

- If $f(\mathbf{x}, \mathbf{A} + \mathbf{B}) = f(\mathbf{x}, \mathbf{A}) + f(\mathbf{x}, \mathbf{B})$, then \mathbf{x} satisfies (11) if and only if \mathbf{x} satisfies (9).
- Under Assumption 1, if \mathbf{x} is feasible in (11), then \mathbf{x} is feasible in (9).

Finally, (11) is tractable due to the following.

Theorem 1 (Bertsimas–Sim): Under Assumption 1, we have the following.

- Constraint (11) is equivalent to

$$f(\mathbf{x}, \mathbf{D}^0) + \Omega \|\mathbf{s}\|^* \leq 0 \quad (12)$$

where

$$s_j = \max\{f(\mathbf{x}, \Delta \mathbf{D}^j), f(\mathbf{x}, -\Delta \mathbf{D}^j)\}.$$

- Equation (12) can be written as $\exists (y, \mathbf{t}) \in \mathbb{R}^{|N|+1}$

$$\begin{aligned} f(\mathbf{x}, \mathbf{D}^0) &\leq -\Omega y \\ f(\mathbf{x}, \Delta \mathbf{D}^j) &\leq t_j \quad \forall j \in N \\ f(\mathbf{x}, -\Delta \mathbf{D}^j) &\leq t_j \quad \forall j \in N \\ \|\mathbf{t}\|^* &\leq y. \end{aligned} \quad (13)$$

Finally, Bertsimas and Sim derive a probability of constraint violation.

Theorem 2 (Bertsimas–Sim): In the model of uncertainty in (10), when we use the l_2 -norm, i.e., $\|\mathbf{s}\|^* = \|\mathbf{s}\|_2$, and under the assumption that $\mathbf{u} \sim \mathcal{N}(\mathbf{0}, I)$, we have the probability bound

$$\mathbb{P}(f(\mathbf{x}, \mathbf{D}) > 0) \leq \frac{\sqrt{e}\Omega}{\alpha} \exp\left(-\frac{\Omega^2}{2\alpha^2}\right)$$

where $\alpha = 1$ for linear programs (LPs), $\alpha = \sqrt{2}$ for SOCPs, and $\alpha = \sqrt{m}$ for SDPs (m is the dimension of the matrix in the SDP).

III. AN EXACT APPROACH USING SDP

In this section, we apply the robust quadratic optimization results to formulate (8) as an SDP. We then show that we can compute optimal solutions to this SDP with a very simple control law.

First, exploiting the linearity of the system, we have the following, straightforward result.

Proposition 2: The cost function (2) for (1) can be written in the form

$$\begin{aligned} J(\mathbf{x}_0, \mathbf{u}, \mathbf{w}) &= 2\mathbf{a}^T \mathbf{x}_0 + \mathbf{x}_0^T \mathbf{A} \mathbf{x}_0 + 2\mathbf{b}^T \mathbf{u} + \mathbf{u}^T \mathbf{B} \mathbf{u} \\ &\quad + 2\mathbf{c}^T \mathbf{w} + \mathbf{w}^T \mathbf{C} \mathbf{w} + 2\mathbf{u}^T \mathbf{D} \mathbf{w} \end{aligned} \quad (14)$$

for appropriate vectors $\mathbf{a} \in \mathbb{R}^{(N \cdot n_x) \times 1}$, $\mathbf{b} \in \mathbb{R}^{(N \cdot n_u) \times 1}$, $\mathbf{c} \in \mathbb{R}^{(N \cdot n_w) \times 1}$ and matrices $\mathbf{A} \in \mathbb{R}^{(N \cdot n_x) \times (N \cdot n_x)}$, $\mathbf{B} \in \mathbb{R}^{(N \cdot n_u) \times (N \cdot n_u)}$, $\mathbf{C} \in \mathbb{R}^{(N \cdot n_w) \times (N \cdot n_w)}$, $\mathbf{D} \in \mathbb{R}^{(N \cdot n_u) \times (N \cdot n_w)}$, and where $\mathbf{B} \succ \mathbf{0}$, $\mathbf{C} \succeq \mathbf{0}$.

Proof: Since the system is linear, we can write the state at any instant k as

$$\mathbf{x}_k = \tilde{\mathbf{A}}_{k-1}\mathbf{x}_0 + \tilde{\mathbf{B}}_{k-1}\mathbf{u} + \tilde{\mathbf{C}}_{k-1}\mathbf{w}$$

where

$$\begin{aligned} \tilde{\mathbf{A}}_{k-1} &= \prod_{i=0}^{k-1} \mathbf{A}_i \\ \tilde{\mathbf{B}}_{k-1} &= \left[\left(\prod_{j=1}^{k-1} \mathbf{A}_j \right) \mathbf{B}_0 \quad \cdots \quad \mathbf{B}_{k-1} \quad \mathbf{0}_{n_x \times (N-k) \cdot n_u} \right] \\ \tilde{\mathbf{C}}_{k-1} &= \left[\left(\prod_{j=1}^{k-1} \mathbf{A}_j \right) \mathbf{C}_0 \quad \cdots \quad \mathbf{C}_{k-1} \quad \mathbf{0}_{n_x \times (N-k) \cdot n_w} \right]. \end{aligned}$$

Now the cost of any state term is written

$$\begin{aligned} &\mathbf{x}_k^T \mathbf{Q}_k \mathbf{x}_k + 2\mathbf{q}_k^T \mathbf{x}_k \\ &= \mathbf{x}_0^T \tilde{\mathbf{A}}_{k-1}^T \mathbf{Q}_k \tilde{\mathbf{A}}_{k-1} \mathbf{x}_0 \\ &\quad + 2\mathbf{x}_0^T \tilde{\mathbf{A}}_{k-1}^T \mathbf{Q}_k \left(\tilde{\mathbf{B}}_{k-1} \mathbf{u} + \tilde{\mathbf{C}}_{k-1} \mathbf{w} \right) \\ &\quad + \mathbf{u}^T \tilde{\mathbf{B}}_{k-1}^T \mathbf{Q}_k \tilde{\mathbf{B}}_{k-1} \mathbf{u} + \mathbf{w}^T \tilde{\mathbf{C}}_{k-1}^T \mathbf{Q}_k \tilde{\mathbf{C}}_{k-1} \mathbf{w} \\ &\quad + 2\mathbf{u}^T \tilde{\mathbf{B}}_{k-1}^T \mathbf{Q}_k \tilde{\mathbf{C}}_{k-1} \mathbf{w} \\ &\quad + 2\mathbf{q}_k^T \left(\tilde{\mathbf{A}}_{k-1} \mathbf{x}_0 + \tilde{\mathbf{B}}_{k-1} \mathbf{u} + \tilde{\mathbf{C}}_{k-1} \mathbf{w} \right). \end{aligned}$$

Thus, the overall cost is clearly written in the form stated above, with

$$\begin{aligned} \mathbf{a} &= \sum_{k=1}^N \tilde{\mathbf{A}}_{k-1}^T \mathbf{q}_k \\ \mathbf{A} &= \sum_{k=1}^N \tilde{\mathbf{A}}_{k-1}^T \mathbf{Q}_k \tilde{\mathbf{A}}_{k-1} \\ \mathbf{b} &= \hat{\mathbf{r}} + \left(\sum_{k=1}^N \tilde{\mathbf{B}}_{k-1}^T \mathbf{Q}_k \tilde{\mathbf{A}}_{k-1} \right) \mathbf{x}_0 + \sum_{k=1}^N \tilde{\mathbf{B}}_{k-1}^T \mathbf{q}_k \\ \mathbf{B} &= \hat{\mathbf{R}} + \sum_{k=1}^N \tilde{\mathbf{B}}_{k-1}^T \mathbf{Q}_k \tilde{\mathbf{B}}_{k-1} \\ \mathbf{c} &= \left(\sum_{k=1}^N \tilde{\mathbf{C}}_{k-1}^T \mathbf{Q}_k \tilde{\mathbf{A}}_{k-1} \right) \mathbf{x}_0 + \sum_{k=1}^N \tilde{\mathbf{C}}_{k-1}^T \mathbf{q}_k \\ \mathbf{C} &= \sum_{k=1}^N \tilde{\mathbf{C}}_{k-1}^T \mathbf{Q}_k \tilde{\mathbf{C}}_{k-1} \\ \mathbf{D} &= \sum_{k=1}^N \tilde{\mathbf{B}}_{k-1}^T \mathbf{Q}_k \tilde{\mathbf{C}}_{k-1} \end{aligned}$$

where

$$\begin{aligned} \hat{\mathbf{r}} &= [\mathbf{r}_0^T \quad \cdots \quad \mathbf{r}_{N-1}^T]^T \\ \hat{\mathbf{R}} &= \text{diag}(\mathbf{R}_0, \dots, \mathbf{R}_{N-1}). \end{aligned}$$

Finally, positive (semi) definiteness of \mathbf{B} and \mathbf{C} follow from positive (semi) definiteness of \mathbf{R}_k and \mathbf{Q}_k . \square

Next, for ease of notation, we will transform the coordinates of the control space.

Proposition 3: To minimize the cost function $J(\mathbf{x}_0, \mathbf{u}, \mathbf{w})$ in Proposition 2 over all $\mathbf{u} \in \mathbb{R}^{N \cdot n_u}$, it is sufficient instead to optimize over all $\mathbf{y} \in \mathbb{R}^{N \cdot n_u}$ the cost function

$$\tilde{J}(\mathbf{x}_0, \mathbf{y}, \mathbf{w}) = \mathbf{y}^T \mathbf{y} + 2\mathbf{h}^T \mathbf{w} + 2\mathbf{y}^T \mathbf{F} \mathbf{w} + \mathbf{w}^T \mathbf{C} \mathbf{w} \quad (15)$$

with $\mathbf{h} = \mathbf{c} - \mathbf{D}^T \mathbf{B}^{-1} \mathbf{b}$, $\mathbf{F} = \mathbf{B}^{-1/2} \mathbf{D}$.

Proof: The proof is immediate from the fact that \mathbf{B}^{-1} exists since $\mathbf{B} \succ \mathbf{0}$, and then using the transformation $\mathbf{u} = \mathbf{B}^{-1/2} \mathbf{y} - \mathbf{B}^{-1} \mathbf{b}$. \square

By Proposition 3, then, (8) is equivalent to the problem

$$\min_{\mathbf{y} \in \mathbb{R}^{N \cdot n_u}} \max_{\mathbf{w} \in \mathcal{W}_\gamma} \tilde{J}(\mathbf{x}_0, \mathbf{y}, \mathbf{w}). \quad (16)$$

This problem may be solved using SDP, as we now show.

Theorem 3: Problem (16) may be solved by the following SDP:

$$\begin{aligned} &\text{minimize } z \\ &\text{subject to} \\ &\quad \begin{bmatrix} \mathbf{I} & \mathbf{y} & \mathbf{F} \\ \mathbf{y}^T & z - \gamma^2 \lambda & -\mathbf{h}^T \\ \mathbf{F}^T & -\mathbf{h} & \lambda \mathbf{I} - \mathbf{C} + \mathbf{F}^T \mathbf{F} \end{bmatrix} \succeq \mathbf{0}, \\ &\quad \lambda \geq 0 \end{aligned} \quad (17)$$

in decision variables \mathbf{y} , z , and λ .

Proof: We first rewrite the problem as

$$\begin{aligned} &\text{minimize } z \\ &\text{subject to } z - \mathbf{y}^T \mathbf{y} - 2\mathbf{h}^T \mathbf{w} - 2\mathbf{y}^T \mathbf{F} \mathbf{w} \\ &\quad - \mathbf{w}^T \mathbf{C} \mathbf{w} \geq 0 \quad \forall \mathbf{w} : \mathbf{w}^T \mathbf{w} \leq \gamma^2. \end{aligned} \quad (18)$$

We may homogenize the system and rewrite this equivalently as

$$\begin{aligned} &\text{minimize } z \\ &\text{subject to } t^2(z - \mathbf{y}^T \mathbf{y}) - 2t\mathbf{h}^T \mathbf{w} - 2t\mathbf{y}^T \mathbf{F} \mathbf{w} \\ &\quad - \mathbf{w}^T \mathbf{C} \mathbf{w} \geq 0 \quad \forall \mathbf{w}, t : \mathbf{w}^T \mathbf{w} \leq \gamma^2 t^2. \end{aligned} \quad (19)$$

Clearly, feasibility of (z, \mathbf{y}) in (19) implies feasibility of (z, \mathbf{y}) in (18) (by setting $t = 1$). For the other direction, assume (z, \mathbf{y}) is feasible in (18) and set $\tilde{\mathbf{w}} = t\mathbf{w}$, where $\mathbf{w}^T \mathbf{w} \leq \gamma^2$. This implies $\tilde{\mathbf{w}}^T \tilde{\mathbf{w}} \leq \gamma^2 t^2$, and

$$\begin{aligned} &2t\mathbf{h}^T \tilde{\mathbf{w}} + 2t\mathbf{y}^T \mathbf{F} \tilde{\mathbf{w}} - \tilde{\mathbf{w}}^T \mathbf{C} \tilde{\mathbf{w}} \\ &= t^2(2\mathbf{h}^T \mathbf{w} + 2\mathbf{y}^T \mathbf{F} \mathbf{w} - \mathbf{w}^T \mathbf{C} \mathbf{w}) \\ &\leq t^2(z - \mathbf{y}^T \mathbf{y}) \end{aligned}$$

where the inequality follows by (18). Thus, the claim is true.

But now we wish to check whether a homogenous quadratic form in (t, \mathbf{w}) is nonnegative over all (t, \mathbf{w}) satisfying another quadratic form. Invoking Lemma 2, we know the constraint holds if and only if there exists a $\lambda \geq 0$ such that

$$\begin{aligned} &\begin{bmatrix} z - \gamma^2 \lambda - \mathbf{y}^T \mathbf{y} & -\mathbf{h}^T - \mathbf{y}^T \mathbf{F} \\ -\mathbf{h} - \mathbf{F}^T \mathbf{y} & \lambda \mathbf{I} - \mathbf{C} \end{bmatrix} \succeq \mathbf{0} \\ &\Updownarrow \\ &\begin{bmatrix} z - \gamma^2 \lambda & -\mathbf{h}^T \\ -\mathbf{h} & \lambda \mathbf{I} - \mathbf{C} + \mathbf{F}^T \mathbf{F} \end{bmatrix} - \begin{bmatrix} \mathbf{y}^T \\ \mathbf{F}^T \end{bmatrix} \begin{bmatrix} \mathbf{y} & \mathbf{F} \end{bmatrix} \succeq \mathbf{0}. \end{aligned}$$

Finally, utilizing Lemma 1 with $B = I$, we see that this is equivalent to

$$\begin{bmatrix} I & \mathbf{y} & \mathbf{F} \\ \mathbf{y}^T & z - \gamma^2 \lambda & -\mathbf{h}^T \\ \mathbf{F}^T & -\mathbf{h} & \lambda I - \mathbf{C} + \mathbf{F}^T \mathbf{F} \end{bmatrix} \succeq \mathbf{0}.$$

Thus we arrive at the desired SDP. \square

There is a tie between the standard DP approach (the Riccati equation) and this SDP, and the connection is not difficult to see.

Corollary 1: With $\gamma = 0$, the optimal solution to SDP (17) solves the Riccati equation, i.e., minimizes the cost-to-go in an expected value sense.

Proof: As argued in Proposition 2, the total cost can be written in the form

$$\begin{aligned} J_{\mathbf{u}}(\mathbf{x}_0) &= 2\mathbf{a}^T \mathbf{x}_0 + \mathbf{x}_0^T \mathbf{A} \mathbf{x}_0 + 2\mathbf{b}^T \mathbf{u} + \mathbf{u}^T \mathbf{B} \mathbf{u} \\ &\quad + 2\mathbf{c}^T \mathbf{w} + \mathbf{w}^T \mathbf{C} \mathbf{w} + 2\mathbf{u}^T \mathbf{D} \mathbf{w}. \end{aligned}$$

With $\gamma = 0$, we require $\mathbf{w} = \mathbf{0} = \mathbb{E}\mathbf{w}$. Hence, we have

$$\max_{\mathbf{w}=\mathbf{0}} J_{\mathbf{u}}(\mathbf{x}_0) = \mathbb{E}[J(\mathbf{x}_0)] + c$$

where $c = -\mathbb{E}[\mathbf{w}^T \mathbf{C} \mathbf{w}]$. Since the goal of the dynamic programming approach is to minimize the expected cost, the equivalence of the two approaches in this case follows. \square

In summary, Theorem 3 provides an exact SDP approach towards solving (16), and in the limiting case $\gamma = 0$, this SDP yields the same solution as the Riccati equation. It is not surprising that the complexity of the problem is that of solving an SDP; in fact, Yao *et al.* [20] have shown that solving the Riccati equation can be cast as an SDP.

A. Simplifying the SDP for Closed-Loop Control

Although Theorem 3 ostensibly provides us with an open-loop policy, we can certainly run this approach in closed loop. We would do this by solving (17) and applying the first n_u components of the solution as the current control. Then, with a new state observation, we update the data to (17) (in fact, only \mathbf{h} depends on the current state, so all other data for the problem can be computed offline) and solve (17) again.

For large problem sizes and applications demanding rapid feedback, however, this approach is clearly impractical. In particular, solving large SDPs of the form of (17) is expensive for large problem sizes, and this a serious drawback in real-time control settings. We would like a simplification that allows us to compute solutions much faster.

In this section, we show that (17) has a very special structure that allows us to dramatically reduce the problem complexity. We will show how to compute optimal policies with only linear operations (i.e., matrix multiplication) in the current state plus a very simple optimization of a convex function over a scalar variable. In short, we will derive for our approach a control law that is an analog of the linear control law (Riccati) in the distributional framework. This control law, while nonlinear in the current state, can nonetheless be computed extremely efficiently and thus may be used very efficiently in closed-loop control.

To begin, we need the following simple observation.

Lemma 3: Consider matrices \mathbf{G}, \mathbf{U} , and \mathbf{V} , of appropriate size, and let $\tilde{\mathbf{U}}$ and $\tilde{\mathbf{V}}$ be the orthogonal complements of \mathbf{U} and \mathbf{V} , respectively (i.e., full-rank matrices such that $\tilde{\mathbf{U}}^T \mathbf{U} = \mathbf{0}$, $\tilde{\mathbf{V}}^T \mathbf{V} = \mathbf{0}$). If there exists a matrix \mathbf{X} such that

$$\mathbf{G} + \mathbf{U} \mathbf{X} \mathbf{V}^T + \mathbf{V} \mathbf{X}^T \mathbf{U}^T \succeq \mathbf{0}$$

then the following hold:

$$\tilde{\mathbf{U}}^T \mathbf{G} \tilde{\mathbf{U}} \succeq \mathbf{0} \quad (20)$$

$$\tilde{\mathbf{V}}^T \mathbf{G} \tilde{\mathbf{V}} \succeq \mathbf{0}. \quad (21)$$

Proof: Clear by multiplying (20) and (21) by $(\tilde{\mathbf{U}}^T, \tilde{\mathbf{U}})$ and $(\tilde{\mathbf{V}}^T, \tilde{\mathbf{V}})$, respectively. \square

Lemma 3 is actually a special case of an ‘‘elimination lemma.’’ The statement is in fact true in both directions when the inequalities are all made strict [7]. We can now apply Lemma 3 to simplify (17). From here on out, we use the notation $\|\mathbf{P}\|_2$ as the spectral norm of a positive semidefinite matrix $\mathbf{P} \in \mathbb{R}^{n \times n}$, i.e.,

$$\|\mathbf{P}\|_2 = \sqrt{\max_{i=1, \dots, n} \lambda_i(\mathbf{P})}$$

where λ_i are the eigenvalues of \mathbf{P} .

Proposition 4: Let z^* be the optimal value of (17). Then $z^* \geq z_R^*$, where z_R^* is the optimal value of the problem

$$\begin{aligned} &\text{minimize} && z \\ &\text{subject to} && \begin{bmatrix} z - \gamma^2 \lambda & -\mathbf{h}^T \\ -\mathbf{h} & \lambda I - \mathbf{C} + \mathbf{F}^T \mathbf{F} \end{bmatrix} \succeq \mathbf{0} \\ &&& \lambda \geq \|\mathbf{C}\|^2 \end{aligned} \quad (22)$$

in decision variables (z, λ) .

Proof: Applying Lemma 3 to the constraint in (17), we see that we can write it in the form

$$\mathbf{G}(z, \lambda) + \mathbf{U} \mathbf{y} \mathbf{V}^T + \mathbf{V} \mathbf{y}^T \mathbf{U}^T \succeq \mathbf{0}$$

where

$$\begin{aligned} \mathbf{G}(z, \lambda) &= \begin{bmatrix} I & \mathbf{0} & \mathbf{F} \\ \mathbf{0}^T & z - \gamma^2 \lambda & -\mathbf{h}^T \\ \mathbf{F}^T & -\mathbf{h} & \lambda I - \mathbf{C} + \mathbf{F}^T \mathbf{F} \end{bmatrix} \\ \mathbf{U} &= \begin{bmatrix} I_{N \cdot n_u} \\ \mathbf{0} \end{bmatrix} \\ \mathbf{V} &= \mathbf{e}_{(N \cdot n_u) + 1}. \end{aligned}$$

We now invoke Lemma 3 with

$$\begin{aligned} \tilde{\mathbf{U}} &= \begin{bmatrix} \mathbf{0} \\ I_{N \cdot n_w + 1} \end{bmatrix} \\ \tilde{\mathbf{V}} &= \begin{bmatrix} I_{N \cdot n_u} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \\ \mathbf{0} & I_{N \cdot n_w} \end{bmatrix} \end{aligned}$$

where all zero matrices are sized appropriately. With these choices, if there exists a \mathbf{y} such that (17) is feasible for a given

(z, λ) , then the following inequalities also hold for such a (z, λ) :

$$\begin{bmatrix} z - \gamma^2 \lambda & -\mathbf{h}^T \\ -\mathbf{h} & \lambda \mathbf{I} - \mathbf{C} + \mathbf{F}^T \mathbf{F} \end{bmatrix} \succeq \mathbf{0}$$

$$\begin{bmatrix} \mathbf{I} & \mathbf{F} \\ \mathbf{F}^T & \lambda \mathbf{I} - \mathbf{C} + \mathbf{F}^T \mathbf{F} \end{bmatrix} \succeq \mathbf{0}$$

and, by Schur complements, the latter inequality is equivalent to $\lambda \mathbf{I} - \mathbf{C} \succeq \mathbf{0} \Rightarrow \lambda \geq \|\mathbf{C}\|^2$, since $\mathbf{C} \succeq \mathbf{0}$. It follows that any (z, λ) that is feasible to (17) is feasible to (22), and hence $z^* \geq z_R^*$. \square

Before analyzing the structure of (22) in more detail, we note the following definition, which we will employ for notational convenience.

Definition 1: Consider a matrix \mathbf{X} such that $\mathbf{X} \succeq \mathbf{0}$ with eigenvalue decomposition written as $\mathbf{X} = \mathbf{Q}\mathbf{\Lambda}\mathbf{Q}^T$. Denote by $\hat{\mathbf{X}}$ the unit full-rank version of \mathbf{X} , with

$$\hat{\mathbf{X}} = \mathbf{Q}\hat{\mathbf{\Lambda}}\mathbf{Q}^T \quad (23)$$

where $\hat{\mathbf{\Lambda}}$ is a diagonal matrix such that

$$[\hat{\mathbf{\Lambda}}]_{ii} = \begin{cases} [\mathbf{\Lambda}]_{ii}, & \text{if } [\mathbf{\Lambda}]_{ii} > 0 \\ 1, & \text{otherwise} \end{cases}$$

Note that $\hat{\mathbf{X}} \succ \mathbf{0}$ always, so $\hat{\mathbf{X}}^{-1}$ always exists; and that if $\mathbf{X} \succ \mathbf{0}$, then $\hat{\mathbf{X}} = \mathbf{X}$.

We now show that (22) can be reduced to a simpler optimization problem involving no semidefinite constraints and just the variable λ . In what follows, we will denote the eigenvalues and eigenvectors of $\mathbf{F}^T \mathbf{F} - \mathbf{C}$ by λ_i and $\mathbf{q}_i, i = 1, \dots, N \cdot n_w$, respectively.

Proposition 5: Problem (22) is equivalent to the convex optimization problem (in single variable λ)

$$\begin{aligned} & \text{minimize} && f(\lambda) \\ & \text{subject to} && \lambda \geq \|\mathbf{C}\|^2 \end{aligned} \quad (24)$$

where

$$f(\lambda) = \begin{cases} \gamma^2 \lambda + \mathbf{h}^T \hat{\mathbf{H}}^{-1}(\lambda) \mathbf{h}, & \text{if } \mathbf{q}_i^T \mathbf{h} = 0 \forall i \in \mathcal{I}(\lambda) \\ +\infty, & \text{otherwise} \end{cases} \quad (25)$$

where $\mathcal{I}(\lambda) = \{i \mid \lambda + \lambda_i = 0\}$ and $\mathbf{H}(\lambda) = \lambda \mathbf{I} - \mathbf{C} + \mathbf{F}^T \mathbf{F}$.

Proof: By Schur complements, a pair (z, λ) is feasible in (22) if and only if

$$\begin{aligned} & \mathbf{H}(\lambda) - \left(\frac{1}{z - \gamma^2 \lambda} \right) \mathbf{h} \mathbf{h}^T \succeq \mathbf{0} \\ \Leftrightarrow & (z - \gamma^2 \lambda) \mathbf{x}^T \mathbf{H}(\lambda) \mathbf{x} \geq (\mathbf{h}^T \mathbf{x})^2 \forall \mathbf{x} \\ \Leftrightarrow & z - \gamma^2 \lambda \geq v^2 \end{aligned}$$

where v is the optimal value of the problem

$$\begin{aligned} & \text{maximize} && \mathbf{h}^T \mathbf{x} \\ & \text{subject to} && \mathbf{x}^T \mathbf{H}(\lambda) \mathbf{x} \leq 1. \end{aligned}$$

Note that feasibility of λ implies that $\lambda + \lambda_i \geq 0$ for all i . Let $\mathcal{I}_+(\lambda) = \{i \mid \lambda + \lambda_i > 0\}$. Then carrying out the above optimization problem, we find

$$\begin{aligned} v &= \begin{cases} \sqrt{\sum_{i \in \mathcal{I}_+(\lambda)} \frac{(\mathbf{q}_i^T \mathbf{h})^2}{\lambda + \lambda_i}}, & \text{if } \mathbf{h}^T \mathbf{q}_j = 0 \forall j \notin \mathcal{I}_+(\lambda) \\ +\infty, & \text{otherwise} \end{cases} \\ &= \begin{cases} \sqrt{\mathbf{h} \hat{\mathbf{H}}^{-1}(\lambda) \mathbf{h}}, & \text{if } \mathbf{h}^T \mathbf{q}_j = 0 \forall j \notin \mathcal{I}_+(\lambda) \\ +\infty, & \text{otherwise.} \end{cases} \end{aligned}$$

From the above equivalences, then, we have (z, λ) feasible to (22) if and only if

$$z \geq \gamma^2 \lambda + \mathbf{h} \hat{\mathbf{H}}^{-1}(\lambda) \mathbf{h}.$$

Since we wish to minimize z , for a fixed λ , we should set z equal to $\gamma^2 \lambda + \mathbf{h} \hat{\mathbf{H}}^{-1}(\lambda) \mathbf{h}$. Thus (22) is equivalent to minimization of $f(\lambda)$ over all $\lambda \geq \|\mathbf{C}\|^2$. \square

We now argue that optimization of $f(\lambda)$ may be done efficiently.

Proposition 6: The function $f(\lambda)$ in (25) satisfies the following.

- $f(\lambda)$ is convex on $\lambda \geq \|\mathbf{C}\|^2$.
- If $\gamma \geq \gamma_{\text{thresh}}$, where

$$\gamma_{\text{thresh}} = \begin{cases} \|\hat{\mathbf{H}}^{-1}(\|\mathbf{C}\|^2) \mathbf{h}\|, & \text{if } \mathbf{q}_i^T \mathbf{h} = 0 \forall i \in \mathcal{I}(\|\mathbf{C}\|^2), \\ +\infty, & \text{otherwise} \end{cases}$$

then $\lambda^* = \|\mathbf{C}\|^2$ minimizes $f(\lambda)$ over all $\lambda \geq \|\mathbf{C}\|^2$.

- If $\gamma > \gamma_{\text{thresh}}$, then the minimizer λ^* of $f(\lambda)$ over all $\lambda \geq \|\mathbf{C}\|^2$ may be found (within tolerance ϵ) in time $\mathcal{O}(\kappa)$, where

$$\kappa = \log_2 \left[\frac{\sqrt{m} \|\mathbf{Q}^T \mathbf{h}\|_\infty - \gamma (\|\mathbf{C}\|^2 + \lambda_{\min})}{\gamma \epsilon} \right] \quad (26)$$

$\lambda_{\min} = \min_{i=1, \dots, N \cdot n_w} \lambda_i$, the columns of \mathbf{Q} are the eigenvectors of $\mathbf{F}^T \mathbf{F} - \mathbf{C}$ and $m \leq N \cdot n_w$ is the number of eigenvectors such that $\mathbf{q}_i^T \mathbf{h} \neq 0$.

Proof:

- We may write $f(\lambda)$ as

$$f(\lambda) = \gamma^2 \lambda + \sum_{i=1}^{N \cdot n_w} \frac{(\mathbf{q}_i^T \mathbf{h})^2}{\lambda + \lambda_i}$$

which is clearly a convex function in λ over $\lambda \geq \|\mathbf{C}\|^2 \Rightarrow \lambda + \lambda_i \geq 0$.

- If $\gamma \geq \gamma_{\text{thresh}}$, we have, for all $\lambda \geq \|\mathbf{C}\|^2$

$$\begin{aligned} f'(\lambda) &= \gamma^2 - \sum_{i=1}^{N \cdot n_w} \frac{(\mathbf{q}_i^T \mathbf{h})^2}{(\lambda + \lambda_i)^2} \\ &\geq \gamma_{\text{thresh}}^2 - \sum_{i=1}^{N \cdot n_w} \frac{(\mathbf{q}_i^T \mathbf{h})^2}{(\lambda + \lambda_i)^2} \\ &= \sum_{i=1}^{N \cdot n_w} \frac{(\mathbf{q}_i^T \mathbf{h})^2}{(\|\mathbf{C}\|^2 + \lambda_i)^2} - \sum_{i=1}^{N \cdot n_w} \frac{(\mathbf{q}_i^T \mathbf{h})^2}{(\lambda + \lambda_i)^2} \\ &\geq 0 \end{aligned}$$

so $f(\lambda)$ is nondecreasing over $\lambda \geq \|\mathbf{C}\|^2$.

c) If $\gamma < \gamma_{\text{thresh}}$, then we must search for $\lambda^* > \|\mathbf{C}\|^2$ such that $f'(\lambda^*) = 0$ [since, by a), $f(\lambda)$ is convex]. This is the same as finding a root $\lambda^* > \|\mathbf{C}\|^2$ of the nonlinear equation

$$\sum_{i=1}^{N \cdot n_w} \frac{(\mathbf{q}_i^T \mathbf{h})^2}{(\lambda + \lambda_i)^2} = \gamma^2. \quad (27)$$

We claim $\lambda^* \leq \bar{\lambda}$, where

$$\bar{\lambda} = \frac{\sqrt{m} \|\mathbf{Q}^T \mathbf{h}\|_\infty}{\gamma} - \min_{i=1, \dots, N \cdot n_w} \lambda_i.$$

Indeed, assume $\lambda^* > \bar{\lambda}$. Then

$$\begin{aligned} f(\lambda^*) &= \sum_{i=1}^{N \cdot n_w} \frac{(\mathbf{q}_i^T \mathbf{h})^2}{(\lambda^* + \lambda_i)^2} \\ &< \sum_{i=1}^{N \cdot n_w} \frac{(\mathbf{q}_i^T \mathbf{h})^2}{(\bar{\lambda} + \lambda_i)^2} \\ &\leq \sum_{i=1}^{N \cdot n_w} \frac{(\mathbf{q}_i^T \mathbf{h})^2}{m \|\mathbf{Q}^T \mathbf{h}\|_\infty^2 / \gamma^2} \\ &= \frac{\gamma^2}{m} \sum_{i=1}^{N \cdot n_w} \frac{(\mathbf{q}_i^T \mathbf{h})^2}{\|\mathbf{Q}^T \mathbf{h}\|_\infty^2} \\ &\leq \gamma^2 \end{aligned}$$

which implies that λ^* cannot be a solution of (27). The result then follows by applying bisection (with tolerance ϵ) on the interval $[\|\mathbf{C}\|^2, \bar{\lambda}]$. \square

We are finally ready for the main result of this section. Given an optimal solution λ^* to (24), we can then compute our closed-loop optimal simply by performing matrix multiplications.

Theorem 4: Let λ^* be the optimal solution to (24). Then the solution $(\lambda^*, z^*, \mathbf{y}^*)$, where

$$\begin{aligned} z^* &= f(\lambda^*) \\ \mathbf{y}^* &= -\mathbf{F} \hat{\mathbf{H}}^{-1}(\lambda^*) \mathbf{h} \end{aligned} \quad (28)$$

is an optimal solution to the SDP (17).

Proof: By Proposition 4, we know that $z^* \geq z_R^*$, and by Proposition 5, we have $z_R^* = f(\lambda^*)$, so if we can just show that the solution $(\lambda^*, z^*, \mathbf{y}^*)$ in (28) is feasible to (17), then we know it must be optimal. First, since $\lambda^* \geq \|\mathbf{C}\|^2 \geq 0$, we have $\lambda^* \geq 0$, as required. For the semidefinite constraint in (17), by Schur complements, we require

$$\begin{bmatrix} z^* - \gamma^2 \lambda^* & & -\mathbf{h}^T \\ & \lambda^* \mathbf{I} - \mathbf{C} + \mathbf{F}^T \mathbf{F} & \\ & & \begin{bmatrix} \mathbf{y}^{*T} \\ \mathbf{F}^T \end{bmatrix} \end{bmatrix} \begin{bmatrix} \mathbf{y}^* \\ \mathbf{F} \end{bmatrix} \succeq \mathbf{0}.$$

As before, let $\mathbf{H}(\lambda^*) = \lambda^* \mathbf{I} - \mathbf{C} + \mathbf{F}^T \mathbf{F}$; note that feasibility of λ^* in (24) implies $\mathbf{H}(\lambda^*) - \mathbf{F}^T \mathbf{F} = \lambda^* \mathbf{I} - \mathbf{C} \succeq \mathbf{0}$, so $\lambda^* \mathbf{I} - \mathbf{C} = \mathbf{G}^T \mathbf{G}$ for some matrix \mathbf{G} . Finally, λ^* minimizes $f(\lambda^*)$ over all $\lambda \geq \|\mathbf{C}\|^2$, so $f(\lambda^*)$ must be finite and thus

$z^* - \gamma^2 \lambda^* = \mathbf{h}^T \hat{\mathbf{H}}^{-1}(\lambda^*) \mathbf{h}$. Putting all of this together, we have

$$\begin{aligned} &\begin{bmatrix} z^* - \gamma^2 \lambda^* & & -\mathbf{h}^T \\ & \lambda^* \mathbf{I} - \mathbf{C} + \mathbf{F}^T \mathbf{F} & \\ & & \begin{bmatrix} \mathbf{y}^{*T} \\ \mathbf{F}^T \end{bmatrix} \end{bmatrix} \begin{bmatrix} \mathbf{y}^* \\ \mathbf{F} \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{h}^T \hat{\mathbf{H}}^{-1}(\lambda^*) \mathbf{h} & & -\mathbf{h}^T \\ & \lambda^* \mathbf{I} - \mathbf{C} + \mathbf{F}^T \mathbf{F} & \\ & & \begin{bmatrix} \mathbf{y}^{*T} \\ \mathbf{F}^T \end{bmatrix} \end{bmatrix} \begin{bmatrix} \mathbf{y}^* \\ \mathbf{F} \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{h}^T \hat{\mathbf{H}}^{-1}(\lambda^*) \mathbf{F}^T \mathbf{F} \hat{\mathbf{H}}^{-1}(\lambda^*) \mathbf{h} & & -\mathbf{h}^T \hat{\mathbf{H}}^{-1}(\lambda^*) \mathbf{F}^T \mathbf{F} \\ & \lambda^* \mathbf{I} - \mathbf{C} + \mathbf{F}^T \mathbf{F} & \\ & & \begin{bmatrix} \mathbf{y}^{*T} \\ \mathbf{F}^T \end{bmatrix} \end{bmatrix} \begin{bmatrix} \mathbf{y}^* \\ \mathbf{F} \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{h}^T \mathbf{P}_1(\lambda^*) \mathbf{h} & & -\mathbf{h}^T \mathbf{P}_2(\lambda^*) \\ & \lambda^* \mathbf{I} - \mathbf{C} & \\ & & \begin{bmatrix} \mathbf{y}^{*T} \\ \mathbf{F}^T \end{bmatrix} \end{bmatrix} \begin{bmatrix} \mathbf{y}^* \\ \mathbf{F} \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{h}^T \hat{\mathbf{H}}^{-1}(\lambda^*) (\lambda^* \mathbf{I} - \mathbf{C}) \hat{\mathbf{H}}^{-1}(\lambda^*) \mathbf{h} & & \mathbf{0} \\ & \lambda^* \mathbf{I} - \mathbf{C} & \\ & & \begin{bmatrix} \mathbf{y}^{*T} \\ \mathbf{F}^T \end{bmatrix} \end{bmatrix} \begin{bmatrix} \mathbf{y}^* \\ \mathbf{F} \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{0} & & -\mathbf{h}^T \hat{\mathbf{H}}^{-1}(\lambda^*) (\lambda^* \mathbf{I} - \mathbf{C}) \\ & \lambda^* \mathbf{I} - \mathbf{C} & \\ & & \begin{bmatrix} \mathbf{y}^{*T} \\ \mathbf{F}^T \end{bmatrix} \end{bmatrix} \begin{bmatrix} \mathbf{y}^* \\ \mathbf{F} \end{bmatrix} \\ &= \begin{bmatrix} -\mathbf{h}^T \hat{\mathbf{H}}^{-1}(\lambda^*) \mathbf{G}^T \\ & \mathbf{G}^T \end{bmatrix} \begin{bmatrix} -\mathbf{G} \hat{\mathbf{H}}^{-1}(\lambda^*) \mathbf{h} \\ \mathbf{G} \end{bmatrix} \\ &\succeq \mathbf{0} \end{aligned}$$

where $\mathbf{P}_1(\lambda^*) = (\hat{\mathbf{H}}^{-1}(\lambda^*) - \hat{\mathbf{H}}^{-1}(\lambda^*) \mathbf{F}^T \mathbf{F} \hat{\mathbf{H}}^{-1}(\lambda^*))$, $\mathbf{P}_2(\lambda^*) = (\mathbf{I} - \hat{\mathbf{H}}^{-1}(\lambda^*) \mathbf{F}^T \mathbf{F})$, and the second-to-last equality follows by factoring $\hat{\mathbf{H}}^{-1}(\lambda^*)$ from these two matrices. This completes the proof. \square

We reiterate that Theorem 4 provides us with a control vector $\mathbf{y} \in \mathbb{R}^{N \cdot n_u}$ for all remaining stages. When we run this in closed loop, however, we would just take the first n_u components and apply that as the control to the current stage.

Thus, we see that the optimal, closed-loop control law for the approach given by (16) may be computed in the following way: first, by computing the optimal value λ^* (which may be done by bisection, via Proposition 6), then by matrix multiplication. The only datum in the SDP (17) that depends on the current state \mathbf{x}_0 is the vector \mathbf{h} . Thus, much of the work may be done offline. We now quantify explicitly the online computational burden.

Corollary 2: Consider the problem setup from (17), with N the number of stages and n_x, n_u , and n_w the sizes of the state, control, and disturbance vectors, respectively. Then the optimal closed-loop policy for a single period may be computed in $\mathcal{O}(N n_w (\kappa + n_x + n_u))$ time, where κ is given in (26).

Proof: Since we only care about the current control, and \mathbf{h} is linear in the current state \mathbf{x}_0 , we may write reexpress the closed-loop control for the current optimal control $\mathbf{u}_{\text{curr}}^*$ as

$$\mathbf{u}_{\text{curr}}^* = \bar{\mathbf{u}} + \tilde{\mathbf{F}} \hat{\mathbf{H}}^{-1}(\lambda^*) \tilde{\mathbf{G}} \mathbf{x}_0$$

where $\tilde{\mathbf{F}} \in \mathbb{R}^{n_u \times N n_w}$, $\tilde{\mathbf{G}} \in \mathbb{R}^{N n_w \times n_x}$, and $\hat{\mathbf{H}}^{-1}(\lambda^*) \in \mathbb{R}^{N n_w \times N n_w}$. Of course, since $H(\lambda) = \lambda \mathbf{I} - \mathbf{C} + \mathbf{F}^T \mathbf{F}$, we may compute an eigenvalue decomposition for $H(0) = \mathbf{Q} \mathbf{\Lambda} \mathbf{Q}^T$ offline, and computing the inverse of $\hat{\mathbf{H}}(\lambda^*)$ may be done simply by adding λ^* to the diagonal elements of $\mathbf{\Lambda}$. Our total computational burden breaks down as follows:

- $\mathcal{O}(N n_w n_x)$ iterations for setting up the optimization of $f(\lambda)$ (i.e., computing $\mathbf{Q}^T \mathbf{h}$);
- $\mathcal{O}(\kappa)$ bisection calls, each requiring a sum over $N n_w$ terms, for a total effort of $\mathcal{O}(\kappa \cdot N n_w)$ searching for λ^* ;

TABLE I
COMPUTATIONAL EFFORT (SECONDS) FOR VARIOUS PROBLEM SIZES ON A
1 GHZ MACHINE. HERE, $\kappa \approx 10$ AND $n_x \approx n_u \approx n_w$ ARE ASSUMED

N	n_w		
	1	10	100
1	1e-8	1e-6	1e-5
10	1e-7	1e-5	1e-4
100	1e-6	1e-4	1e-3
1000	1e-5	1e-3	1e-2
10000	1e-4	1e-2	1e-1

- right-matrix multiplication of the current state: $\mathcal{O}(Nn_w n_x)$;
- scaling the resulting vector componentwise by $\lambda^* + \lambda_i$: $\mathcal{O}(Nn_w)$;
- left-matrix multiplication to obtain optimal control: $\mathcal{O}(Nn_w n_u)$.

The total effort required is thus

$$\mathcal{O}(Nn_w(1 + \kappa + n_x + 1 + n_u)) = \mathcal{O}(Nn_w(\kappa + n_x + n_u)).$$

□

We note that the computational effort, for all other inputs fixed, grows *linearly* with the number of stages N . Of course, by computing the matrices offline, we are reducing the computational effort by increasing storage requirements. Our total memory usage is to store N matrices of sizes $n_u \cdot Nn_w$ and $Nn_w \cdot n_x$, for a total memory requirement of $\mathcal{O}(N^2 n_w (n_x + n_u))$. Hence, memory increases quadratically with the time horizon N .

Table I illustrates order-of-magnitude estimates for the computational time for various problem sizes.

IV. AN INNER APPROXIMATION USING SOCP

In the presence of constraints, we cannot use the simplification results of Section III-A. This means for constrained control with feedback, we would need to solve a problem of the same form as (17) with constraints at each stage. For large problems and applications demanding fast decisions, this will not be feasible. This motivates us to find a simplification of the exact SDP in (17).

Here we will develop an inner approximation using SOCP [i.e., any feasible solution to the SOCP will be feasible to SDP (17)]. We will exploit the robust conic optimization results highlighted in Section II-D.

Recall that our cost-to-go can be written as

$$\tilde{J}(x_0, \mathbf{y}, \mathbf{w}) = \mathbf{y}^T \mathbf{y} + 2(\mathbf{h} + \mathbf{F}^T \mathbf{y})^T \mathbf{w} + \mathbf{w}^T \mathbf{C} \mathbf{w}.$$

As before, we would like to find the policy \mathbf{y} , which minimizes the maximum value of $\tilde{J}(x_0, \mathbf{y}, \mathbf{w})$ over all $\mathbf{w} \in \mathbb{R}^{N \cdot n_w}$ in some ellipsoidal uncertainty set. We may write this problem as

$$\begin{aligned} & \text{minimize } z \\ & \text{subject to } f(\mathbf{y}, \mathbf{w}) \leq z - \|\mathbf{y}\|_2^2 \quad \forall \mathbf{w} \in \mathcal{W}_\Omega \end{aligned} \quad (29)$$

where $f(\mathbf{y}, \mathbf{w}) = 2(\mathbf{h} + \mathbf{F}^T \mathbf{y})^T \mathbf{w} + \mathbf{w}^T \mathbf{C} \mathbf{w}$. Our uncertainty model is

$$\mathcal{W}_\Omega = \left\{ \sum_{j=1}^{N \cdot n_w} u_j \mathbf{e}^j, \left\| \mathbf{u} \right\|_2 \leq \Omega \right\} \quad (30)$$

where $\{\mathbf{e}^1, \dots, \mathbf{e}^{N \cdot n_w}\}$ is any orthonormal basis of $\mathbb{R}^{N \cdot n_w}$. Note that \mathcal{W}_Ω is precisely the same uncertainty set utilized in Section III, with Ω assuming an analogous role as γ . In the framework of Section II-D, we are using the assignments $\mathbf{D} = \mathbf{w}$, $\mathbf{D}^0 = \mathbf{0}$, and $\mathbf{D}^j = \mathbf{e}^j$.

From here, we would like to directly apply the results from Section II-D. The difficulty, however, is that the quadratic term $\mathbf{w}^T \mathbf{C} \mathbf{w}$ causes $f(\mathbf{y}, \mathbf{w})$ to violate Assumption 1b). We remove this difficulty with a slight relaxation of (29).

Proposition 7: Consider the problem

$$\begin{aligned} & \text{minimize } z \\ & \text{subject to } \tilde{f}(\mathbf{y}, \mathbf{w}) \leq z - \|\mathbf{y}\|_2^2 - \Omega^2 \|\mathbf{C}\|_2^2 \quad \forall \mathbf{w} \in \mathcal{W}_\Omega \end{aligned} \quad (31)$$

where $\tilde{f}(\mathbf{y}, \mathbf{w}) = 2(\mathbf{h} + \mathbf{F}^T \mathbf{y})^T \mathbf{w}$ and \mathcal{W}_Ω is as described in (30). We then have the following.

- If a solution (z, \mathbf{y}) is feasible in (31), then it is also feasible in (29).
- The function $\tilde{f}(\mathbf{y}, \mathbf{w})$ satisfies Assumption 1.

Proof:

- Let (z, \mathbf{y}) be feasible in (31), and consider any $\mathbf{w} \in \mathcal{W}_\Omega$. Then we have

$$\begin{aligned} f(\mathbf{y}, \mathbf{w}) &= 2(\mathbf{h} + \mathbf{F}^T \mathbf{y})^T \mathbf{w} + \mathbf{w}^T \mathbf{C} \mathbf{w} \\ &\leq 2(\mathbf{h} + \mathbf{F}^T \mathbf{y})^T \mathbf{w} + \Omega^2 \|\mathbf{C}\|_2^2 \end{aligned} \quad (32)$$

$$\leq z - \|\mathbf{y}\|_2^2 \quad (33)$$

where (32) follows from the fact that

$$\max_{\|\mathbf{x}\| \leq \Omega} \mathbf{x}^T \mathbf{C} \mathbf{x} = \Omega^2 \|\mathbf{C}\|_2^2$$

and (33) follows from feasibility of (z, \mathbf{y}) in (31). Thus, (z, \mathbf{y}) is also feasible in the original formulation in (29).

- This follows trivially, since $\tilde{f}(\mathbf{y}, \mathbf{w})$ is linear in \mathbf{w} . □
Now that we have cast the problem in the framework of Section II-D, we may apply the corresponding results. This leads us to our formulation of (29) as a second-order cone problem, as we now illustrate.

Theorem 5: Consider the SOCP

$$\begin{aligned} & \text{minimize } z \\ & \text{subject to } \|\mathbf{y}\|_2^2 \leq z - \Omega \tilde{\gamma} - \Omega^2 \|\mathbf{C}\|_2^2 \\ & \quad |2(\mathbf{h} + \mathbf{F}^T \mathbf{y})^T \mathbf{e}^j| \leq t_j, \\ & \quad j = 1, \dots, N \cdot n_w, \quad \|\mathbf{t}\|_2 \leq \tilde{\gamma} \end{aligned} \quad (34)$$

in decision variables $(z, \mathbf{y}, \hat{\mathbf{y}}, \mathbf{t})$. If (z, \mathbf{y}) are part of a feasible solution to (34), then (z, \mathbf{y}) are also feasible in (29).

Proof: From Theorem 1, (34) is equivalent to

$$\begin{aligned} & \text{minimize } z \\ & \text{subject to } \tilde{f}(\mathbf{y}, \mathbf{0}) + \max_{(\mathbf{u}, \mathbf{v}) \in \mathcal{V}_\Omega} \sum_{j=1}^{N \cdot n_w} \{ \tilde{f}(\mathbf{y}, \mathbf{e}^j) u_j + \tilde{f}(\mathbf{y}, -\mathbf{e}^j) v_j \} \\ & \leq z - \|\mathbf{y}\|_2^2 - \Omega^2 \|\mathbf{C}\|_2^2 \end{aligned}$$

where

$$\mathcal{V}_\Omega = \left\{ (\mathbf{u}, \mathbf{v}) \in \mathbb{R}_+^{2(N \cdot n_w)} \mid \|\mathbf{u} + \mathbf{v}\|_2 \leq \Omega \right\}.$$

Now, since $\tilde{f}(\mathbf{y}, \mathbf{w})$ is linear in \mathbf{w} , this problem is equivalent to (31), by Proposition 1a). Finally, invoking Proposition 7a), we have that feasibility of (z, \mathbf{y}) in (31) implies feasibility in (29), and thus we are done. \square

Theorem 5 thus gives us an inner approximation to the exact problem given in (29). This is in contrast to Theorem 3, which solves the problem exactly using SDP (and, in the unconstrained case, can be simplified via the results in Section III-A). Theorem 5 gives us an SOCP formulation, which is a significant reduction in complexity from the SDP. In addition, we expect this approximation to be quite tight, as the *only* inequality we have exploited is $\mathbf{w}^T \mathbf{C} \mathbf{w} \leq \Omega^2 \|\mathbf{C}\|_2^2$, which holds for all $\mathbf{w} \in \mathcal{W}_\Omega$. We now quantify this difference.

Corollary 3: Let z_{SDP}^* and z_{SOCP}^* be the optimal values of the SDP and SOCP (with $\Omega = \gamma$) given in Theorems 3 and 5, respectively. Then we have

$$z_{\text{SOCP}}^* - z_{\text{SDP}}^* \leq 2\gamma \|\mathbf{h}\|_2. \tag{35}$$

Proof: Note that we have

$$\begin{aligned} z_{\text{SOCP}}^* &= \gamma^2 \|\mathbf{C}\|_2^2 + \min_{\mathbf{y}} \left[\|\mathbf{y}\|_2^2 + 2 \max_{\mathbf{w}: \|\mathbf{w}\|_2 \leq \gamma} (\mathbf{h} + \mathbf{F}^T \mathbf{y})^T \mathbf{w} \right] \\ &\leq \gamma^2 \|\mathbf{C}\|_2^2 + 2 \max_{\mathbf{w}: \|\mathbf{w}\|_2 \leq \gamma} \mathbf{h}^T \mathbf{w} \\ &= \gamma^2 \|\mathbf{C}\|_2^2 + 2\gamma \|\mathbf{h}\|_2 \end{aligned}$$

where the inequality follows from feasibility of $\mathbf{y} = \mathbf{0}$. On the other hand, positive semidefiniteness of the matrix in Theorem 3 requires

$$\begin{aligned} z_{\text{SDP}}^* &\geq \|\mathbf{y}\|_2^2 + \gamma^2 \lambda \\ &\geq \gamma^2 \lambda \\ &\geq \gamma^2 \|\mathbf{C}\|_2^2 \end{aligned}$$

where the last line follows from Lemma 2. The result in (35) now follows.

V. CONSTRAINTS AND PERFORMANCE GUARANTEES

We now demonstrate the modelling power of the approaches developed in Theorems 3 and 5. In particular, we show that both approaches readily lend themselves towards handling a wide variety of constraints. These constraints fit into three categories:

control constraints, probabilistic guarantees on the state, and probabilistic guarantees on the cost function. For the probabilistic guarantees, we will assume the disturbances are independently and normally distributed.

The model for uncertainty proposed in Section III is deterministic and relies on a norm-bounded disturbance vector. While it is true that this model does not seem, at first glance, to apply to random variables which are unbounded, the purpose of the probability results within this section is to show that the optimal solutions based our uncertainty model *do* in fact have reasonable performance guarantees even when the underlying disturbance vectors obey a different uncertainty model, namely, one admitting a probabilistic description. In particular, we utilize the normal distribution because of its analytical convenience and prevalence throughout much of the control literature. Thus, even though our original model utilizes a bounded uncertainty model, the resulting solutions still perform well under a stochastic model with unbounded disturbance vectors. This is the underlying thrust for the probability results of this section.

We present the results here for both the SDP and SOCP frameworks. Since the presence of constraints destroys the simple control law for the SDP from Section III-A, however, the SOCP is more viable in a constrained, closed-loop control setting (in fact, we reiterate that this was the primary motivation for the development of the SOCP approach).

Throughout this section, we will make claims about the ‘‘complexity type’’ of the problem being unchanged. By this, we mean the SDP remains an SDP and the SOCP remains an SOCP. We implicitly appeal to the fact that the class of SDP problems includes the class of SOCP problems, and thus we may add second-order cone constraints to an SDP without increasing its complexity type.

We turn first to the simplest case of control constraints.

A. Control Constraints

We will show that both approaches may handle any convex quadratic constraint on the control vector. In this and the following section, we temporarily revert to the traditional notation \mathbf{u} for the controls and note that the simple affine transformation listed in Proposition 3 allows us to implement these constraints in our \mathbf{y} control space. We first need the following well-known result.

Proposition 8: The quadratic constraint $\mathbf{x}^T \mathbf{x} \leq t$ is equivalent to the second-order cone constraint

$$\left\| \begin{pmatrix} \mathbf{x} \\ \frac{t-1}{2} \end{pmatrix} \right\|_2 \leq \frac{t+1}{2}.$$

Proof: This is a standard result (see, e.g., [3]) that follows by noting that $t = ((t+1)^2)/(4) - ((t-1)^2)/(4)$. \square

We now have the rather straightforward result of this section.

Theorem 6: Any control constraints of the form

$$\|\mathbf{G}\mathbf{u}\|_2^2 + 2\mathbf{g}^T \mathbf{u} + \hat{g} \leq 0 \tag{36}$$

where $\mathbf{G} \in \mathbb{R}^{(N \cdot n_u) \times (N \cdot n_u)}$, $\mathbf{g} \in \mathbb{R}^{(N \cdot n_u) \times 1}$, and $\hat{g} \in \mathbb{R}$ may be suitably added to (17) and (34) without increasing their respective complexity types.

Proof: By Proposition 8, we may write (36) as

$$\left\| \begin{pmatrix} \mathbf{G}u \\ -\mathbf{g}^T \mathbf{u} - \frac{(\hat{g}+1)}{2} \end{pmatrix} \right\|_2 \leq -\mathbf{g}^T \mathbf{u} - \frac{(\hat{g}-1)}{2}$$

which is a second-order cone constraint and hence may be added to either problem without raising the complexity type. \square

Note that Theorem 6 implies we can tractably deal with any polyhedral or ellipsoidal constraints on the control.

B. Probabilistic State Guarantees

Since the state \mathbf{x} of the system is not exactly known, any constraints on \mathbf{x} can only be enforced in a probabilistic sense. To ensure probabilistic guarantees in what follows, we will operate under the typical assumption that the disturbances \mathbf{w} are independently and normally distributed.³

Assumption 2: The disturbances \mathbf{w} are independently and normally distributed with zero mean, i.e.,

$$\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}).$$

Note that, if instead we have $\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \Sigma)$, where $\Sigma \succ \mathbf{0}$, we may simply rotate coordinates and multiply the \mathbf{C}_k matrices in the dynamics in (1) accordingly.

We now show how to explicitly ensure that linear constraints on the state will hold with a desirably high probability. The notations Φ and χ_n stand for the cumulative distribution functions of standard normal and n -degree chi-squared variables, respectively, and the notation $\|\cdot\|_{\mathbf{A}}$ will represent the Euclidean norm induced under the matrix $\mathbf{A} \succeq \mathbf{0}$, i.e., $\|\mathbf{x}\|_{\mathbf{A}} = \sqrt{\mathbf{x}^T \mathbf{A} \mathbf{x}}$.

Theorem 7: Consider a linear system described by (1), with the state written as

$$\mathbf{x} = \bar{\mathbf{A}}\mathbf{x}_0 + \bar{\mathbf{B}}\mathbf{u} + \bar{\mathbf{C}}\mathbf{w}$$

for appropriate matrices $\bar{\mathbf{A}}$, $\bar{\mathbf{B}}$, and $\bar{\mathbf{C}}$. Then under Assumption 2, we have the following.

a) The constraint

$$\mathbf{p}^T \mathbf{u} \leq q \quad (37)$$

where $\mathbf{p} = \bar{\mathbf{B}}^T \mathbf{g}$, $q = \hat{g} - \mathbf{g}^T \bar{\mathbf{A}}\mathbf{x}_0 - \|\bar{\mathbf{C}}^T \mathbf{g}\|_2 \Phi^{-1}(1 - \epsilon)$ implies the following guarantee:

$$\mathbb{P}(\mathbf{g}^T \mathbf{x} > \hat{g}) \leq \epsilon. \quad (38)$$

b) The constraint

$$\|\bar{\mathbf{A}}\mathbf{x}_0 + \bar{\mathbf{B}}\mathbf{u}\|_{\mathbf{H}} \leq r \quad (39)$$

where $r = 1 - \|\bar{\mathbf{C}}^T \mathbf{H} \bar{\mathbf{C}}\|_2 \chi_{Nk}^{-1}(1 - \epsilon)$ implies the following guarantee:

$$\mathbb{P}(\mathbf{x}^T \mathbf{H} \mathbf{x} > 1) \leq \epsilon \quad (40)$$

where $\mathbf{H} \succeq \mathbf{0}$.

³Of course, other distributional assumptions may be made; we present the normality assumption primarily because a) it provides the cleanest analytical results and b) it is the most common assumption in the literature.

Proof:

a) We have

$$\begin{aligned} \mathbb{P}(\mathbf{g}^T \mathbf{x} > \hat{g}) &= \mathbb{P}(\mathbf{g}^T \bar{\mathbf{C}}\mathbf{w} > \hat{g} - \mathbf{g}^T \bar{\mathbf{A}}\mathbf{x}_0 - \mathbf{g}^T \bar{\mathbf{B}}\mathbf{u}) \\ &= 1 - \Phi \left(\frac{\hat{g} - \mathbf{g}^T \bar{\mathbf{A}}\mathbf{x}_0 - \mathbf{g}^T \bar{\mathbf{B}}\mathbf{u}}{\|\bar{\mathbf{C}}^T \mathbf{g}\|_2} \right) \end{aligned}$$

and the result follows by setting this less than or equal to ϵ and inverting Φ .

b) We have

$$\begin{aligned} \mathbb{P}(\mathbf{x}^T \mathbf{H} \mathbf{x} > 1) &= \mathbb{P}(\|\bar{\mathbf{A}}\mathbf{x}_0 + \bar{\mathbf{B}}\mathbf{u} + \bar{\mathbf{C}}\mathbf{w}\|_{\mathbf{H}} > 1) \\ &\leq \mathbb{P}(\|\bar{\mathbf{A}}\mathbf{x}_0 + \bar{\mathbf{B}}\mathbf{u}\|_{\mathbf{H}} + \|\bar{\mathbf{C}}\mathbf{w}\|_{\mathbf{H}} > 1) \\ &\leq \mathbb{P}(\{\|\bar{\mathbf{C}}^T \mathbf{H} \bar{\mathbf{C}}\|_2 \|\mathbf{w}\|_2 > 1 - \|\bar{\mathbf{A}}\mathbf{x}_0 + \bar{\mathbf{B}}\mathbf{u}\|_{\mathbf{H}}\}) \\ &= 1 - \chi_{Nk} \left(\frac{1 - \|\bar{\mathbf{A}}\mathbf{x}_0 + \bar{\mathbf{B}}\mathbf{u}\|_{\mathbf{H}}}{\|\bar{\mathbf{C}}^T \mathbf{H} \bar{\mathbf{C}}\|_2} \right) \end{aligned}$$

and, again, the result follows by setting this less than or equal to ϵ and inverting χ_{Nk} . \square

Both parts of Theorem 7 are constraints that can be added to either the SDP or SOCP approaches without increasing their respective complexity types. Note that the constraint in part a) of the theorem is exact, while the constraint in b) is somewhat conservative. Care must be taken to ensure that \mathbf{H} and ϵ do not result in a constraint that forces the problem to be infeasible.

C. Probabilistic Performance Guarantees

In this section, we analyze the probability distribution of the cost-to-go function. We first derive a bound on the performance distribution of the cost-to-go function for a given control policy \mathbf{y} under Assumption 2, then describe the protection guarantees and expected losses for both (17) and (34). Finally, we show how to probabilistically ensure certain levels of performance.

We emphasize that the results proven here in terms of performance guarantees are for *open-loop* control. In general, analyzing our approach in a feedback context seems difficult. Instead, we will study the closed-loop performance computationally in the Section VII.

For a given policy \mathbf{y} , the cost function is a random variable (a function of the random disturbances)

$$\tilde{J}(\mathbf{x}_0, \mathbf{y}, \mathbf{w}) = \mathbf{w}^T \mathbf{C} \mathbf{w} + 2(\mathbf{h} + \mathbf{F}^T \mathbf{y})^T \mathbf{w} + \mathbf{y}^T \mathbf{y}.$$

For our results in this section, we need a slightly stronger assumption.

Assumption 3: In addition to Assumption 2, we have $\mathbf{C} \succ \mathbf{0}$. We see that under Assumption 3, we have

$$\mathbb{E}[\tilde{J}(\mathbf{x}_0, \mathbf{y}, \mathbf{w})] = \|\mathbf{y}\|_2^2 + \text{trace}(\mathbf{C})$$

with $\text{Tr}(\mathbf{C}) = \sum_{i=1}^{N \cdot n_w} \lambda_i$, where λ_i are the eigenvalues of \mathbf{C} . We now derive a key result; the proof is quite similar to a proof from [6].

Proposition 9: Under Assumption 3, we have

$$\begin{aligned} & \mathbb{P}(\tilde{J}(\mathbf{x}_0, \mathbf{y}, \mathbf{w}) > z) \\ & \leq c_\rho \exp\left(-\frac{1}{2\text{Tr}(\mathbf{C})} \left[z - \|\mathbf{y}\|_2^2 - \frac{1}{\rho-1} \|\mathbf{g}\|_{\mathbf{C}^{-1}}^2 \right]\right) \end{aligned} \quad (41)$$

where

$$c_\rho = \left(\frac{\rho}{\rho-1}\right)^{\rho/2} \quad (42)$$

$$\rho = \frac{\text{Tr}(\mathbf{C})}{\max_{i=1, \dots, N \cdot n_w} \lambda_i} > 1 \quad (43)$$

$$\mathbf{g} = \mathbf{h} + \mathbf{F}^T \mathbf{y}. \quad (44)$$

Proof: Let $\lambda_i > 0, i = 1, \dots, N \cdot n_w$, be the eigenvalues of $\mathbf{C} \succ \mathbf{0}$, and $\mathbf{Q}^T \mathbf{\Lambda} \mathbf{Q}$ be its eigenvalue decomposition. We then have

$$\begin{aligned} & \mathbb{P}(\tilde{J}(\mathbf{x}_0, \mathbf{y}, \mathbf{w}) > z) \\ & = \mathbb{P}(\mathbf{w}^T \mathbf{C} \mathbf{w} + 2(\mathbf{h} + \mathbf{F}^T \mathbf{y})^T \mathbf{w} > z - \|\mathbf{y}\|_2^2) \\ & = \mathbb{P}\left(\sum_{i=1}^{N \cdot n_w} \lambda_i (v_i + f_i/\lambda_i)^2 > z - \|\mathbf{y}\|_2^2 + \|\mathbf{g}\|_{\mathbf{C}^{-1}}^2\right) \\ & = \mathbb{P}\left(\sum_{i=1}^{N \cdot n_w} \lambda_i u_i^2 > \tilde{z}\right) \end{aligned}$$

where we employ the transformations $\mathbf{v} = \mathbf{Q} \mathbf{w}$, $\mathbf{f} = \mathbf{Q}(\mathbf{h} + \mathbf{F}^T \mathbf{y})$ and we have $u_i \sim \mathcal{N}(f_i/\lambda_i, 1)$, independent. Finally, $\tilde{z} = z - \|\mathbf{y}\|_2^2 + \|\mathbf{g}\|_{\mathbf{C}^{-1}}^2$ for notational convenience. Continuing, we have

$$\begin{aligned} \mathbb{P}\left(\sum_{i=1}^{N \cdot n_w} \lambda_i u_i^2 > \tilde{z}\right) & \leq \frac{\mathbb{E}\left[\exp\left(\theta \sum_{i=1}^{N \cdot n_w} \lambda_i u_i^2\right)\right]}{\exp(\theta \tilde{z})} \\ & = \frac{\prod_{i=1}^{N \cdot n_w} \mathbb{E}\left[\exp\left(\theta \lambda_i u_i^2\right)\right]}{\exp(\theta \tilde{z})} \\ & = \frac{\prod_{i=1}^{N \cdot n_w} \mathbb{E}\left[\exp\left(\frac{u_i^2}{\beta}\right)^{\beta \theta \lambda_i}\right]}{\exp(\theta \tilde{z})} \\ & \leq \frac{\prod_{i=1}^{N \cdot n_w} \mathbb{E}\left[\exp\left(\frac{u_i^2}{\beta}\right)^{\beta \theta \lambda_i}\right]}{\exp(\theta \tilde{z})} \end{aligned}$$

where we require $\theta > 0, \beta \theta \lambda_i \leq 1$, and $\beta > 2$. The first line above follows from the Markov inequality, the second follows from independence of the u_i , and the last line follows from Jensen's inequality. Now noting that under $u_i \sim \mathcal{N}(f_i/\lambda_i, 1)$, we have

$$\mathbb{E}\left[\exp\left(\frac{u_i^2}{\beta}\right)\right] = \sqrt{\frac{\beta}{\beta-2}} \exp\left(\frac{1}{\beta-2} (f_i/\lambda_i)^2\right)$$

we have, after some rearranging with

$$\theta = 1/(\beta \max_{i=1, \dots, N \cdot n_w} \lambda_i)$$

$$\begin{aligned} & \frac{\prod_{i=1}^{N \cdot n_w} \mathbb{E}\left[\exp\left(\frac{u_i^2}{\beta}\right)^{\beta \theta \lambda_i}\right]}{\exp(\theta \tilde{z})} \\ & = \frac{\prod_{i=1}^{N \cdot n_w} \left(\sqrt{\frac{\beta}{\beta-2}} \exp\left(\frac{1}{\beta-2} (f_i/\lambda_i)^2\right)\right)^{\beta \theta \lambda_i}}{\exp(\theta \tilde{z})} \\ & = \frac{\left(\frac{\beta}{\beta-2}\right)^{\rho/2} \exp\left(\frac{2\rho}{\beta(\beta-2)} \frac{\|\mathbf{g}\|_{\mathbf{C}^{-1}}^2}{\text{Tr} \mathbf{C}}\right)}{\exp\left(\frac{\rho}{\beta} \left[\frac{z - \|\mathbf{y}\|_2^2}{\text{Tr} \mathbf{C}}\right]\right)}. \end{aligned}$$

Finally, we set $\beta = 2\rho > 2$ and the result follows. \square

Note that the bound in (41) is for *any* policy \mathbf{y} , and we have in no way imposed the structure of either of the approaches developed in Sections III or IV. Now utilizing the structure of the SDP and the SOCP, we may quantify more precisely what we gain in terms of performance protection with both approaches. This protection comes at the price of some degradation of expected performance, and we also quantify this decrease.

Theorem 8: Under Assumption 3, and with $\mathbb{E}[\tilde{J}(\mathbf{x}_0, \mathbf{0}, \mathbf{w})]$ the expected value of the Riccati approach, we have the following.

- a) If $(\mathbf{y}_{\text{SDP}}^*, z_{\text{SDP}}^*)$ are feasible in SDP (17), then the expected performance loss is bounded as

$$\mathbb{E}[\tilde{J}(\mathbf{x}_0, \mathbf{y}_{\text{SDP}}^*, \mathbf{w})] - \mathbb{E}[\tilde{J}(\mathbf{x}_0, \mathbf{0}, \mathbf{w})] \leq 2\|\mathbf{h}\|_{2\gamma} \quad (45)$$

while we gain the following level of probabilistic protection:

$$\mathbb{P}(\tilde{J}(\mathbf{x}_0, \mathbf{y}_{\text{SDP}}^*, \mathbf{w}) > z_{\text{SDP}}^*) \leq \alpha \exp\left(-\frac{\rho}{2} \gamma^2\right) \quad (46)$$

where

$$\begin{aligned} \alpha & = c_\rho \exp\left(\frac{1}{2(\rho-1)\text{Tr} \mathbf{C}} \delta^2\right) \\ \delta & = \|\mathbf{C}^{-1}\|_2 \left[\|\mathbf{h}\|_2 + \|\mathbf{F} \mathbf{F}^T\|_2 \sqrt{2\|\mathbf{h}\|_{2\gamma}}\right]. \end{aligned}$$

- b) If $(\mathbf{y}_{\text{SOCP}}^*, z_{\text{SOCP}}^*)$ are feasible in SOCP (34), then the expected performance loss is bounded as

$$\mathbb{E}[\tilde{J}(\mathbf{x}_0, \mathbf{y}_{\text{SOCP}}^*, \mathbf{w})] - \mathbb{E}[\tilde{J}(\mathbf{x}_0, \mathbf{0}, \mathbf{w})] \leq 4\|\mathbf{h}\|_{2\Omega} \quad (47)$$

while we gain the following level of probabilistic protection:

$$\mathbb{P}(\tilde{J}(\mathbf{x}_0, \mathbf{y}_{\text{SOCP}}^*, \mathbf{w}) > z_{\text{SOCP}}^*) \leq \sqrt{\frac{e}{2}} \Omega \exp\left(-\frac{\Omega^2}{4}\right). \quad (48)$$

Proof:

- a) For the expected loss from the Riccati approach, note that the optimal Riccati solution, from Corollary 1, is $\mathbf{y} = \mathbf{0}$, and that the expected performance of *any* policy is just given by

$$\mathbb{E}[\tilde{J}(\mathbf{x}_0, \mathbf{y}, \mathbf{w})] = \|\mathbf{y}\|_2^2 + \text{Tr}\mathbf{C}.$$

Therefore, noting that feasibility of $(\mathbf{y}_{\text{SDP}}^*, z_{\text{SDP}}^*)$ in SDP (17) requires $\|\mathbf{y}_{\text{SDP}}^*\|_2^2 \leq z_{\text{SDP}}^* - \|\mathbf{C}\|_2^2 \gamma^2$, we have

$$\begin{aligned} & \mathbb{E}[\tilde{J}(\mathbf{x}_0, \mathbf{y}_{\text{SDP}}^*, \mathbf{w})] - \mathbb{E}[\tilde{J}(\mathbf{x}_0, \mathbf{0}, \mathbf{w})] \\ &= \|\mathbf{y}_{\text{SDP}}^*\|_2^2 \\ &\leq z_{\text{SDP}}^* - \|\mathbf{C}\|_2^2 \gamma^2 \\ &= \min_{\mathbf{y}} \left[\|\mathbf{y}\|_2^2 + \max_{\mathbf{w} \in \mathcal{W}_\gamma} \mathbf{w}^T \mathbf{C} \mathbf{w} \right. \\ &\quad \left. + 2(\mathbf{h} + \mathbf{F}^T \mathbf{y})^T \mathbf{w} \right] - \|\mathbf{C}\|_2^2 \gamma^2 \\ &\leq \max_{\mathbf{w} \in \mathcal{W}_\gamma} [\mathbf{w}^T \mathbf{C} \mathbf{w} + 2(\mathbf{h} + \mathbf{F}^T \mathbf{y})^T \mathbf{w}] - \|\mathbf{C}\|_2^2 \gamma^2 \\ &\leq \|\mathbf{C}\|_2^2 \gamma^2 + 2\|\mathbf{h}\|_2 \gamma - \|\mathbf{C}\|_2^2 \gamma^2 \\ &= 2\|\mathbf{h}\|_2 \gamma. \end{aligned}$$

The first line follows from feasibility of $(\mathbf{y}_{\text{SDP}}^*, z_{\text{SDP}}^*)$ as stated above; the next line follows from the definition of z_{SDP}^* ; the next line follows by noting that $\mathbf{y} = \mathbf{0}$ is feasible in the SDP; and the second-to-last line follows by bounding the maximum value of the given function over all $\mathbf{w} \in \mathcal{W}_\gamma$.

For the probabilistic guarantee of (46), we note, from Proposition 9 and feasibility of $(\mathbf{y}_{\text{SDP}}^*, z_{\text{SDP}}^*)$ that

$$\begin{aligned} & \mathbb{P}(\tilde{J}(\mathbf{x}_0, \mathbf{y}_{\text{SDP}}^*, \mathbf{w}) > z_{\text{SDP}}^*) \\ &\leq c_\rho \exp\left(-\frac{1}{2\text{Tr}(\mathbf{C})} [z_{\text{SDP}}^* \right. \\ &\quad \left. - \|\mathbf{y}_{\text{SDP}}^*\|_2^2 - \frac{1}{\rho-1} \|\mathbf{g}\|_{\mathbf{C}^{-1}}^2]\right) \\ &\leq c_\rho \exp\left(-\frac{1}{2\text{Tr}(\mathbf{C})} \left[\|\mathbf{C}\|_2^2 \gamma^2 - \frac{1}{\rho-1} \|\mathbf{g}\|_{\mathbf{C}^{-1}}^2\right]\right) \\ &= c_\rho \exp\left(\frac{1}{2(\rho-1)\text{Tr}(\mathbf{C})} \|\mathbf{g}\|_{\mathbf{C}^{-1}}^2\right) \exp\left(-\frac{\rho}{2} \gamma^2\right). \end{aligned}$$

Now, we have

$$\begin{aligned} \|\mathbf{g}\|_{\mathbf{C}^{-1}} &= \|\mathbf{h} + \mathbf{F}^T \mathbf{y}_{\text{SDP}}^*\|_{\mathbf{C}^{-1}} \\ &\leq \|\mathbf{C}^{-1}\|_2 \|\mathbf{h} + \mathbf{F}^T \mathbf{y}_{\text{SDP}}^*\|_2 \\ &\leq \|\mathbf{C}^{-1}\|_2 [\|\mathbf{h}\|_2 + \|\mathbf{F}^T \mathbf{y}_{\text{SDP}}^*\|_2] \\ &\leq \|\mathbf{C}^{-1}\|_2 [\|\mathbf{h}\|_2 + \|\mathbf{F}\mathbf{F}^T\|_2 \|\mathbf{y}_{\text{SDP}}^*\|_2] \\ &\leq \|\mathbf{C}^{-1}\|_2 [\|\mathbf{h}\|_2 + \|\mathbf{F}\mathbf{F}^T\|_2 \sqrt{2}\|\mathbf{h}\|_2 \gamma] \end{aligned}$$

where we have repeatedly used matrix norm bounds and the Schwartz inequality, and, in the last line, utilized the bound for $\|\mathbf{y}_{\text{SDP}}^*\|_2^2$ derived in the expected loss above. The result now follows.

- b) The expected loss for the SOCP follows exactly analogously to the proof for the expected loss for the SDP in a) by noting that feasibility of $(\mathbf{y}_{\text{SOCP}}^*, z_{\text{SOCP}}^*)$ in the SOCP implies feasibility in the SDP (Proposition 7), replacing γ with Ω , and then noting the result of Corollary 3 (namely,

$z_{\text{SOCP}}^* - z_{\text{SDP}}^* \leq 2\Omega\|\mathbf{h}\|_2$). The probabilistic bound follows by directly applying Theorem 2. \square

Theorem 8 quantifies the expected loss and a probabilistic protection level for both approaches. Note that the expected loss from the Riccati equation can be bounded by a quantity linear in the size of the uncertainty set (γ or Ω). Moreover, the protection level bounds are both of a similar nature ($\mathcal{O}(\gamma \exp(-\gamma^2))$ and $\mathcal{O}(\Omega \exp(-\Omega^2))$).

In addition to simply describing performance, we may also want to explicitly protect against certain threshold performance levels. We now show how to do this.

Theorem 9: Under Assumption 3, the convex quadratic constraint

$$\mathbf{y}^T \mathbf{P} \mathbf{y} + 2\mathbf{q}^T \mathbf{y} \leq r \quad (49)$$

where

$$\mathbf{P} = \mathbf{I} + \frac{1}{\rho-1} \mathbf{F} \mathbf{C}^{-1} \mathbf{F}^T \succ \mathbf{0} \quad (50)$$

$$\mathbf{q} = \frac{1}{\rho-1} \mathbf{F} \mathbf{C}^{-1} \mathbf{h} \quad (51)$$

$$r = z + 2\text{Tr}\mathbf{C} \ln(\epsilon/c_\rho) - \frac{1}{\rho-1} \mathbf{h}^T (\mathbf{C})^{-1} \mathbf{h} \quad (52)$$

implies the following guarantee:

$$\mathbb{P}(\tilde{J}(\mathbf{x}_0, \mathbf{y}, \mathbf{w}) > z) \leq \epsilon. \quad (53)$$

Proof: The proof follows directly from Proposition 9 in straightforward fashion, noting the implications

$$\begin{aligned} & \mathbf{y}^T \mathbf{P} \mathbf{y} + 2\mathbf{q}^T \mathbf{y} \leq r \\ &\Rightarrow \|\mathbf{y}\|_2^2 + \frac{1}{\rho-1} \|\mathbf{g}\|_{\mathbf{C}^{-1}}^2 \leq z + 2\text{Tr}(\mathbf{C}) \ln(\epsilon/c_\rho) \\ &\Rightarrow c_\rho \exp\left(-\frac{1}{2\text{Tr}(\mathbf{C})} \left[z - \|\mathbf{y}\|_2^2 - \frac{1}{\rho-1} \|\mathbf{g}\|_{\mathbf{C}^{-1}}^2\right]\right) \leq \epsilon \\ &\Rightarrow \mathbb{P}(\tilde{J}(\mathbf{x}_0, \mathbf{y}, \mathbf{w}) > z) \leq \epsilon \end{aligned}$$

where \mathbf{P} , \mathbf{q} , and r are defined in (50)–(52). Positive definiteness of \mathbf{P} follows from positive definiteness of \mathbf{I} and \mathbf{C} , and thus the constraint is a convex quadratic constraint. \square

We see that (49) is a convex quadratic constraint and hence may be added (in the same manner as in the proof of Theorem 6) to either approach without increasing their respective complexity types. Note that we may only ensure against appreciably high levels of cost. In fact, a simple necessary (but not sufficient) condition to retain feasibility of the problem is the requirement

$$z \geq -2\text{Tr}(\mathbf{C}) \ln(\epsilon/c_\rho).$$

VI. IMPERFECT STATE INFORMATION

In some cases, we may not know the state \mathbf{x}_0 of the system exactly. Rather, we may only have an estimate $\hat{\mathbf{x}}_0$ of the current state. In standard dynamic programming texts (see [4]), it is shown that in the case where noise-corrupted state observations of the form

$$\mathbf{v}_k = \mathbf{H}_k \mathbf{x}_k + \delta_k$$

are available, with \mathbf{H}_k known matrices and $\boldsymbol{\delta}_k$ additive noise with finite second moment, then the resulting optimal policy is a modified Riccati equation. Here, we assume the following model for the state estimate $\hat{\mathbf{x}}_0$:

$$\hat{\mathbf{x}}_0 = \mathbf{x}_0 + \boldsymbol{\eta} \tag{54}$$

where $\boldsymbol{\eta}$ is a noise term with some distribution. We will show that the form of the cost-to-go function is unchanged by the added uncertainty in the state by now viewing the disturbances as $\hat{\mathbf{w}} = [\boldsymbol{\eta}^T \ \mathbf{w}^T]^T$. As a consequence, we can apply either of the robust approaches to the problem with imperfect state information of the form given in (54).

Proposition 10: With noisy estimates of the state given by (54), the cost-to-go can be written in the form

$$\begin{aligned} \hat{J}_{\mathbf{u}}(\hat{\mathbf{x}}_0) = & \hat{\mathbf{x}}_0^T \hat{\mathbf{A}} \hat{\mathbf{x}}_0 + 2\hat{\mathbf{a}}^T \hat{\mathbf{x}}_0 + 2\hat{\mathbf{b}}^T \mathbf{u} + \mathbf{u}^T \mathbf{B} \mathbf{u} \\ & + 2\hat{\mathbf{w}}^T \hat{\mathbf{c}} + \hat{\mathbf{w}}^T \hat{\mathbf{C}} \hat{\mathbf{w}} + 2\mathbf{u}^T \hat{\mathbf{D}} \hat{\mathbf{w}} \end{aligned} \tag{55}$$

where $\hat{\mathbf{A}}$, $\hat{\mathbf{a}}$, and \mathbf{B} are as in Proposition 2, $\hat{\mathbf{b}}$ is as in Proposition 2 with $\hat{\mathbf{x}}_0$ replacing \mathbf{x}_0 , and

$$\begin{aligned} \hat{\mathbf{c}} &= \left[\begin{array}{c} -\hat{\mathbf{a}} - \hat{\mathbf{A}}\hat{\mathbf{x}}_0 \\ \left(\sum_{k=0}^{N-1} \tilde{\mathbf{C}}_{k-1}^T \mathbf{Q}_k \tilde{\mathbf{A}}_{k-1} \right) \hat{\mathbf{x}}_0 + \sum_{k=0}^{N-1} \tilde{\mathbf{C}}_{k-1}^T \mathbf{q}_k \end{array} \right] \\ \hat{\mathbf{C}} &= \left[\begin{array}{cc} \hat{\mathbf{A}} & -\sum_{k=1}^N \tilde{\mathbf{C}}_{k-1}^T \mathbf{Q}_k \tilde{\mathbf{C}}_{k-1} \\ -\sum_{k=1}^N \tilde{\mathbf{C}}_{k-1}^T \mathbf{Q}_k \tilde{\mathbf{C}}_{k-1} & \mathbf{C} \end{array} \right] \\ \hat{\mathbf{D}} &= \left[\begin{array}{c} -\sum_{k=1}^N \left(\mathbf{R}_{k+1} + \tilde{\mathbf{B}}_{k-1}^T \mathbf{Q}_k \tilde{\mathbf{B}}_{k-1} \right) \\ D \end{array} \right] \end{aligned}$$

Furthermore, the matrix $\hat{\mathbf{C}}$ is positive semidefinite.

Proof: The proof follows by recalling the original, perfect state cost form of

$$\begin{aligned} J_{\mathbf{u}}(\mathbf{x}_0) = & 2\hat{\mathbf{a}}^T \mathbf{x}_0 + \mathbf{x}_0^T \hat{\mathbf{A}} \mathbf{x}_0 + 2\hat{\mathbf{b}}^T \mathbf{u} + \mathbf{u}^T \mathbf{B} \mathbf{u} + 2\mathbf{c}^T \mathbf{w} \\ & + \mathbf{w}^T \mathbf{C} \mathbf{w} + 2\mathbf{u}^T \mathbf{D} \mathbf{w} \end{aligned}$$

(14). Simply substituting $\mathbf{x}_0 = \hat{\mathbf{x}}_0 - \boldsymbol{\eta}$ and collecting terms, (55) follows. To see that $\hat{\mathbf{C}} \succeq \mathbf{0}$, note, using the original definitions of $\hat{\mathbf{A}}$ and \mathbf{C} , we have

$$\begin{aligned} \hat{\mathbf{C}} &= \left[\begin{array}{cc} \sum_{k=1}^N \tilde{\mathbf{A}}_{k-1}^T \mathbf{Q}_k \tilde{\mathbf{A}}_{k-1} & -\sum_{k=1}^N \tilde{\mathbf{C}}_{k-1}^T \mathbf{Q}_k \tilde{\mathbf{C}}_{k-1} \\ -\sum_{k=1}^N \tilde{\mathbf{C}}_{k-1}^T \mathbf{Q}_k \tilde{\mathbf{C}}_{k-1} & \sum_{k=1}^N \tilde{\mathbf{C}}_{k-1}^T \mathbf{Q}_k \tilde{\mathbf{C}}_{k-1} \end{array} \right] \\ &= \sum_{k=1}^N \left[\begin{array}{cc} \tilde{\mathbf{A}}_{k-1}^T \mathbf{Q}_k \tilde{\mathbf{A}}_{k-1} & -\tilde{\mathbf{C}}_{k-1}^T \mathbf{Q}_k \tilde{\mathbf{C}}_{k-1} \\ -\tilde{\mathbf{C}}_{k-1}^T \mathbf{Q}_k \tilde{\mathbf{C}}_{k-1} & \tilde{\mathbf{C}}_{k-1}^T \mathbf{Q}_k \tilde{\mathbf{C}}_{k-1} \end{array} \right] \\ &= \sum_{k=1}^N [\tilde{\mathbf{A}}_{k-1} \quad -\tilde{\mathbf{C}}_{k-1}]^T \mathbf{Q}_k [\tilde{\mathbf{A}}_{k-1} \quad -\tilde{\mathbf{C}}_{k-1}] \\ &\succeq \mathbf{0} \end{aligned}$$

where the last line follows, since it is a sum of similarity transformations of positive definite matrices ($\mathbf{Q}_k \succ \mathbf{0}$). \square

VII. COMPUTATIONAL RESULTS

We have written routines for solving the control law of Section III-A as well as SOCP (34) in Section IV. The routines

TABLE II
AVERAGE RELATIVE COST INCREASE [SDP (UPPER), SOCP (MIDDLE)] AND STABILITY INCREASE [SDP (LOWER)] ALL VERSUS RICCATI FOR VARIOUS γ AND DISTURBANCE DISTRIBUTIONS

Average relative cost increase (%), SDP						
γ	.001	.01	.1	1	10	
σ	.01	.00	.05	4.0	51	51
	.1	.00	.13	4.8	54	54
	1	.04	.45	5.0	48	49
	10	.00	.02	.20	2.8	19
Average relative cost increase (%), SOCP						
γ	.001	.01	.1	1	10	
σ	.01	.00	.05	4.0	135	140
	.1	.00	.13	4.8	102	131
	1	.04	.45	5.0	145	150
	10	.00	.02	.20	4.6	48
Stability increase (%), SDP						
γ	.001	.01	.1	1	10	
σ	.01	.16	1.5	17	110	120
	.10	.13	1.5	15	80	86
	1	.06	.57	5.1	15	19
	10	.00	.01	.13	.41	7.7

have been implemented in a Matlab environment and the SeDuMi [16] package has been used for the underlying optimization problems. In this section, we explore computationally the performance of our approach in closed-loop control in a variety of ways.

A. Performance in the Unconstrained Case

Here we compare the performance of the optimal policy in Theorem 4 to that given by the Riccati equation for a problem without constraints. We considered a simple, ten-stage problem with time-invariant state, control, and disturbance matrices $\mathbf{A} = 1$, $\mathbf{B} = 1$, and $\mathbf{C} = 1$, initial state $\mathbf{x}_0 = -1$. The cost function was given by $\mathbf{Q}_k = (1/2)^k$, $\mathbf{R}_k = (1/2)^k$, and $\mathbf{q}_k = \mathbf{r}_k = 0$ for all $k \in \{0, \dots, 10\}$.

We ran 1000 trials of the closed-loop policies for the Riccati approach and the control law of Theorem 4 and tabulated the average percentage increase in cost (over Riccati) for various values of γ . Disturbances vectors were generated at each iteration by $\mathcal{N}(0, \sigma^2)$, where σ is a parameter we varied. Table II lists the results.

We observe the following from these computational results.

- 1) For γ small, this approach does not result in a marked increase in expected cost.
- 2) For γ beyond a certain value, the expected cost increase does not change. This is not surprising, since for γ large enough, we have $\gamma > \gamma_{\text{thresh}}$ [Proposition 6b]) at each iteration. In this case, there will be no change in the policy given by Theorem 4 for further increases in γ , so the performance, on average, will not change.
- 3) The policies given by Theorem 4 are more conservative than the traditional, Riccati approach. As such, the distribution in the cost is more stable for larger γ . Here, we measure stability in terms of the standard deviation of the first upper tail moment. Specifically, if we denote the cost distribution under a control policy π by the random variable J_π , then the stability results reported are

$$\text{Stability}_\pi = \sigma(\mathbb{E}[\max(J_\pi - \mathbb{E}J_\pi, 0)]).$$

We chose the upper tail moment because downside variability in the cost is potentially beneficial. The policies given by Theorem 4 give, for the most part, significantly more stable policies than the Riccati policy. Thus, by varying γ , there is a tradeoff between expected value of the cost and *variability* of the cost.

- 4) Although we only report results for a $N = 10$ stage problem here, the results are similar for other dimensions and other problem instances.

B. SOCP Performance in the Unconstrained Case

The SOCP approach from Section IV has been developed for use in the constrained case (i.e., when we cannot use the control law given in Theorem 4). As it is an approximation to SDP (17), however, it is relevant to examine how good this approximation is when we have no constraints. We ran 100 trials of the problem instance from the previous section and compared the performance of SOCP (34) to the control law in Theorem 4, with both approaches run closed-loop. The results are shown in Table II, again versus γ and under various disturbance distributions.

Note that the shape of the expected cost increase for SOCP (34) versus γ is essentially the same as that for the control law of Theorem 4, just with a higher asymptote for large γ . For $\gamma \ll 1$, however, the performance of the SOCP approximation is essentially indistinguishable from that of the control law in Theorem 4.

C. Effect of Constraints on Runtime

Here we present resulting run times for various values of N for a problem with $n_x = n_u = n_w = 1$, $\mathbf{A}_k = \mathbf{B}_k = \mathbf{1}$, $\mathbf{C}_k = 1/(2N)$, $\mathbf{Q}_k = \mathbf{1}$, $\mathbf{q}_k = \mathbf{0}$, $\mathbf{R}_k = \mathbf{0}$, $\mathbf{r}_k = \mathbf{0}$, and $\mathbf{x}_0 = \mathbf{1}$. For all trials we use $\Omega = 0$. The machine is running Linux on a 2.2 GHz processor with 1.0 GB RAM. The results are listed in Table III. Note that in this case we are solving the problem in open loop. We note the following.

- 1) The presence of any of the constraints listed does not result in marked increases in run time for fixed N .
- 2) For the objective guarantee, we use $\epsilon = 0.2$ for $N = 10$ and $\epsilon = 0.1$ for all other N . In the unconstrained case, we can only guarantee that $\mathbb{P}(J_{\mathbf{u}_0}(\mathbf{x}_0) > 2E[J_{\mathbf{u}_0}(\mathbf{x}_0)]) \leq 1$, where \mathbf{u}_0 is the optimal policy in that case. As we can see, however, without significant increase in expected cost we are able to ensure that this bad event occurs with probability no greater than ϵ .
- 3) The constraints $\mathbf{u} \geq 0$ result in very large increases in expected cost. This is merely due to the fact that it is a very restrictive restraint and not a drawback of the proposed approaches.
- 4) Although we do not report the run-times here, we ran this simulation for the SDP with the listed constraints as well. Typically this runs $\mathcal{O}(N)$ longer than the SOCP, solidifying our assertion that the SOCP is much more suitable to efficient, closed-loop control.

D. Performance on a Problem With Constraints

Here we compared the performance of SOCP (34) versus the optimal policy for a five-stage problem with the constraints $u_k \geq 0$ for $k \in \{0, \dots, 4\}$. The problem data were $\mathbf{A}_k =$

TABLE III
RUN-TIME IN SECONDS AND COST INCREASE FROM UNCONSTRAINED FOR THE SOCP APPROACH WITH VARIOUS CONSTRAINTS

Run Time (sec.)			
N	10	100	1000
No constraints	0.06	0.47	346.00
$\mathbb{P}(J_{\mathbf{u}_0}(\mathbf{x}_0) > 2E[J_{\mathbf{u}_0}(\mathbf{x}_0)]) \leq \epsilon$	0.16	1.87	938.10
$\mathbb{P}(x_i < 0) \leq .01$	0.14	1.54	1279.90
$\mathbf{u} \geq 0$	0.11	1.57	1550.50
$\mathbb{P}(x_i < 0) \leq .01$ & $\mathbf{u} \geq 0$	0.11	1.55	1546.90
Increase from Unconstrained, Expected Cost Per Stage (%)			
N	10	100	1000
$\mathbb{P}(J_{\mathbf{u}_0}(\mathbf{x}_0) > 2E[J_{\mathbf{u}_0}(\mathbf{x}_0)]) \leq \epsilon$	0.79	1.35	0.00
$\mathbb{P}(x_i < 0) \leq .01$	86.24	91.89	100.00
$\mathbf{u} \geq 0$	7261.24	6.82E5	7.15E7
$\mathbb{P}(x_i < 0) \leq .01$ & $\mathbf{u} \geq 0$	7261.24	6.82E5	7.15E7

TABLE IV
AVERAGE RELATIVE COST INCREASE FOR SOCP VERSUS RICCATI FOR VARIOUS γ AND DISCOUNT FACTORS

Average relative cost increase (%), SOCP						
γ	.001	.01	.1	1	10	
β	.1	.00	.00	.02	1.5	21
	.5	.00	.00	.57	16	24
	1.0	-.10	-1.5	.00	13	37

$\mathbf{1}$, $\mathbf{B}_k = \mathbf{1}$, and $\mathbf{C}_k = \mathbf{1}$ for $k = \{0, \dots, 4\}$, initial state $\mathbf{x}_0 = -.6$. The cost function was given by $\mathbf{Q}_k = \beta^k$, $\mathbf{R}_k = \beta^k$, and $\mathbf{q}_k = \mathbf{0}$, $\mathbf{r}_k = -\beta^k$, for all $k \in \{0, \dots, 4\}$. Here, β is a discounting factor that we varied in the simulations. Note that smaller β implies that the cost associated with later stages is less important. We compared this to the optimal control law (computed with enumeration over a grid approximation) when the disturbances satisfy $w_k \sim \mathcal{N}(0, 0.1)$. The results are illustrated in Table IV, and we note the following.

- 1) The form of the increase in expected relative cost is very similar to the unconstrained case. In particular, for $\gamma \ll 1$, there is very little degradation in performance from the optimal policy. For large enough γ , the performance approaches an upper limit (in the range 20–40% in this case) and does not go beyond that.
- 2) The performance loss is better for smaller β . This makes intuitive sense, as our approach does not exactly capture the tradeoff between current costs and future costs. When β is smaller, then, future costs are less relevant and SOCP (34), which is myopic, will perform better.
- 3) For $\beta = 1$, we see that our approach appears to outperform the optimal policy for small γ . This can be attributed due to small error in computing the optimal policy due to the discrete grid approximation.

VIII. CONCLUSION

A primary open question of interest is how to simplify this approach even further in the presence of constraints. The best we have here is to solve an SOCP at each step. From Section VII, we see how the complexity of this problem grows faster than linearly with the size of the problem. For large enough problems, solving this in closed loop will become overly burdensome. It remains to be seen if the constrained SOCP has a structure that can be simplified for more efficient computations.

ACKNOWLEDGMENT

The authors would like to thank Prof. L. El Ghaoui for providing valuable insights and useful references that helped them develop this paper. They also thank the Associate Editor and the reviewers for a number of helpful comments.

REFERENCES

- [1] J. C. M. Baños and M. Papageorgiou, "A linear programming approach to large-scale linear optimal control problems," *IEEE Trans. Autom. Control*, vol. 40, pp. 971–977, May 1995.
- [2] T. Başar and P. Bernhard, *H[∞]-Optimal Control and Related Minimax Design Problems: A Dynamic Game Approach*. Boston, MA: Birkhäuser, 1995.
- [3] A. Ben-Tal and A. Nemirovski, "Lectures on modern convex optimization," *MPS-SIAM Series on Optimization*, 2001.
- [4] D. P. Bertsekas, *Dynamic Programming and Optimal Control*, 2nd ed. Boston, MA: Athena Scientific, 2000.
- [5] D. Bertsimas and A. W. Lo, "Optimal control of execution costs," *J. Financial Markets*, pp. 1–50, 1998.
- [6] D. Bertsimas and M. Sim, "Tractable approximations of robust conic optimization problems," *Math. Program. B*, to be published.
- [7] S. Boyd, L. El-Ghaoui, E. Feron, and V. Balakrishnan, "Linear matrix inequalities in system and control theory," *SIAM Studies Appl. Math.*, 1994.
- [8] X. Chen and X. Zhou, "Stochastic LQ control with conic control constraints on an infinite time horizon," *SIAM J. Contr. Optim.*, vol. 43, pp. 1120–1150, 2004.
- [9] L. El-Ghaoui and H. Lebret, "Robust solutions to least-squares problems with uncertain data," *SIAM J. Matrix Anal. Applicat.*, Oct. 1997.
- [10] A. Hansson, "A primal-dual interior-point method for robust optimal control of linear discrete-time systems," *IEEE Trans. Autom. Control*, vol. 45, pp. 1639–1655, Sep. 2000.
- [11] F. Paganini, *A Set-Based Methodology for White Noise Modeling*, unpublished.
- [12] P. L. D. Peres and L. C. Geromel, "An alternate numerical solution to the LQP," *IEEE Trans. Autom. Control*, vol. 39, pp. 198–202, Jan. 1994.
- [13] M. A. Rami and L. El Ghaoui, "LMI optimization for nonstandard Riccati equations arising in stochastic control," *IEEE Trans. Autom. Control*, Jul. 1997.
- [14] M. A. Rami and X. Y. Zhou, "Linear matrix inequalities, riccati equations, and indefinite stochastic linear controls," *IEEE Trans. Autom. Control*, vol. 45, pp. 1131–1143, Jun. 2000.
- [15] A. Rantzer and M. Johansson, "Piecewise linear quadratic optimal control," *IEEE Trans. Autom. Control*, vol. 45, pp. 629–637, Apr. 2000.

- [16] J. F. Sturm, SeDuMi Matlab Optimization Toolbox [Online]. Available: <http://fewcal.kub.nl/sturm/software/sedumi.html>
- [17] E. I. Verriest and G. A. Pajunen, "Quadratically saturated regulator for constrained linear systems," *IEEE Trans. Autom. Control*, vol. 41, pp. 992–995, Jul. 1996.
- [18] P. G. Voulgaris, "Optimal \mathcal{H}_2/l_1 control via duality theory," *IEEE Trans. Autom. Control*, vol. 40, pp. 1881–1888, Nov. 1995.
- [19] J. C. Willems, "Least squares stationary optimal control and the algebraic Riccati equation," *IEEE Trans. Autom. Control*, vol. AC-16, pp. 621–634, Dec. 1971.
- [20] D. D. Yao, S. Zhang, and X. Y. Zhou, "Stochastic linear-quadratic control via semidefinite programming," *SIAM J. Contr. Optim.*, vol. 40, no. 3, pp. 801–823.
- [21] K. Zhou, J. C. Doyle, and K. Glover, *Robust Optimal Control*. Englewood Cliffs, NJ: Prentice-Hall, 1996.



Dimitris Bertsimas received the B.S. degree in electrical engineering and computer science from the National Technical University of Athens, Greece, in 1985, the M.S. degree in operations research from the Massachusetts Institute of Technology (MIT), Cambridge, in 1987, and the Ph.D. degree in applied mathematics and operations research from MIT in 1988.

He is the Boeing Professor of Operations Research and the co-Director of the Operations Research Center, Massachusetts Institute of Technology. Since 1988, he has been with MIT's Sloan School of Management. His research interests include optimization, stochastic systems, data mining, and their application. He has published widely and is a former Area Editor of *Operations Research*.

Prof. Bertsimas is a member of the National Academy of Engineering. He has received numerous research awards, including the Erlang prize (1996), the SIAM prize in optimization (1996), the Bodossaki prize (1998), and the Presidential Young Investigator award (1991–1996).



David B. Brown received the B.S. and M.S. degrees in electrical engineering from Stanford University, Stanford, CA, in 2000 and 2001, respectively, and the Ph.D. degree in electrical engineering and computer science from the Massachusetts Institute of Technology (MIT), Cambridge, in 2006.

He is currently an Assistant Professor of decision sciences in the Fuqua School of Business, Duke University, Durham, NC. His research interests include optimization under uncertainty, robust optimization, connections to risk theory, and financial applications.