

Design of Affine Controllers via Convex Optimization

Joëlle Skaf, *Member, IEEE*, and Stephen P. Boyd, *Fellow, IEEE*

Abstract—We consider a discrete-time time-varying linear dynamical system, perturbed by process noise, with linear noise corrupted measurements, over a finite horizon. We address the problem of designing a general affine causal controller, in which the control input is an affine function of all previous measurements, in order to minimize a convex objective, in either a stochastic or worst-case setting. This controller design problem is not convex in its natural form, but can be transformed to an equivalent convex optimization problem by a nonlinear change of variables, which allows us to efficiently solve the problem. Our method is related to the classical Q -design procedure for time-invariant, infinite-horizon linear controller design, and the more recent purified output control method. We illustrate the method with applications to supply chain optimization and dynamic portfolio optimization, and show the method can be combined with model predictive control techniques when perfect state information is available.

Index Terms—Affine controller, dynamical system, dynamic linear programming (DLP), linear exponential quadratic Gaussian (LEQG), linear quadratic Gaussian (LQG), model predictive control (MPC), proportional-integral-derivative (PID).

I. INTRODUCTION

WE consider a discrete-time time-varying linear dynamical system operating over a finite horizon, perturbed by process noise, with linear noise corrupted measurements of the state, and a cost function that is convex in the state and input trajectories. We will consider two models of the process and measurement noises: stochastic, with the objective measured as the mean cost; and worst-case, in which the disturbances lie in a given set, and the objective is the maximum cost, over all possible disturbances. We require that the input or action must be causal, *i.e.*, a function (or policy) of the previously observed measurements. With a stochastic noise model, the problem of finding a policy that minimizes the objective is the general stochastic control problem, also called stochastic optimization with recourse [3], [14], [18], [57], [75]. With a worst-case noise model, it is called robust optimization with recourse or robust dynamic programming [9], [11], [12], [78].

Even with linear dynamics and convex underlying cost function, the stochastic control problem is very difficult to solve. In the special case when noise-free measurements of the state are

available, the optimal control policy is a function of the current state, which can, in principle, be determined by solving the dynamic programming or Bellman equation to find the cost-to-go function [4], [5]. While this gives us much insight, it does not lead to a practical method, except when the state dimension is very small (say, no more than 3). Another special case in which we obtain a practical solution is when the cost is quadratic and the disturbances are Gaussian, *i.e.*, the linear quadratic Gaussian (LQG) problem. In this special case, the optimal control policy is affine, *i.e.*, a linear function of the past measurements, plus a constant [2], [63]. Another case in which the optimal control policy is known, and affine, is when the cost is the exponential of a quadratic function, and the disturbances are Gaussian, which is the linear exponential quadratic Gaussian (LEQG) or risk-sensitive LQG problem [75, vol. 1, §19].

Many methods have been proposed to find suboptimal policies that can work well in practice. Perhaps the simplest methods are those from classical linear feedback control techniques, such as proportional-integral-derivative (PID) control [35]. One very effective technique that can be used when a noiseless measurement of the state is available is model predictive control (MPC) [6], [50], [53], [75], which goes by many other names, including dynamic matrix control [29], rolling horizon planning [28], and dynamic linear programming (DLP) [66]. MPC is based on solving a convex optimization problem at each step, with the unknown future disturbances replaced with estimates available at the current time (such as conditional means); but only the current action or input is used. At the next step, the same problem is solved, this time using the exact value of the current state, which is now known from the measurement. Another approach goes under the name approximate dynamic programming [13], [14], [55], in which some estimate of the optimal value function, or optimal policy, is found. While many of the methods assume that perfect state measurements are available, they can be used with an estimate of the state, obtained by an observer or state estimator [19]. (The optimal controller in the very special case of LQG can be expressed as an optimal state feedback, applied to an optimal estimate of the state; but in essentially all other cases, such a separation principle does not hold.) Of course, there are many custom methods for specific problems.

In this paper we describe another method that yields a suboptimal controller, that, like many of the others listed above, can work well in practice. We limit our search to *affine controllers*, in which the control policy functions are all affine. Such controllers are characterized by a constant portion, and a set of controller gains that relate current and past measurements to the current input or action; in particular, they are parametrized by a finite number of real parameters. While affine policies are known to be optimal in the LQG case, they are certainly not optimal in the general case. In the language of stochastic programming with recourse, affine controllers are very specialized

Manuscript received April 30, 2008; revised September 26, 2008; accepted March 11, 2010. First published March 22, 2010; current version published November 03, 2010. This work was supported by the NSF award ECS-0423905, by NSF award 0529426, by DARPA award N66001-06-C-2021, by NASA award NNX07AE11A, by AFOSR award FA9550-06-1-0514, and by AFOSR award FA9550-06-1-0312. Recommended by Associate Editor P. A. Parrilo.

The authors are with the Information Systems Lab, Electrical Engineering Department, Stanford University, Stanford, CA 94305-9510 USA (e-mail: jskaf@stanford.edu; jskaf@google.com; boyd@stanford.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TAC.2010.2046053

and limited; the only recourse we have is linear in the measurements. While there is limited theoretical basis to expect that such controllers will work well (e.g., see [15] for proof of optimality of affine controllers in the context of one-dimensional, constrained, multi-stage robust optimization), examples suggest that in many cases they work quite well. This should not be too surprising, since PID control (a very special case of affine controllers) is widely used in industrial control.

In any case, the point of this paper is not to endorse or recommend affine controllers, especially over the other suboptimal techniques described above. The main point of this paper is to show that *the globally optimal affine controller can be effectively computed by convex optimization*. More accurately, we reduce the problem to solving a convex stochastic or worst-case robust optimization problem. In some cases such problems can be solved exactly; in many others, an approximation method, such as sampling, must be used. These approximations can be made arbitrarily tight (at the cost of computational effort), and often yield very good results with nominal effort.

Computing the optimal affine controller (coefficients) requires solving a single convex stochastic or worst-case robust optimization problem *offline*, so the time required to solve the problem is not critical. When the affine controller is running online, the computational load is extremely small, and consists of evaluating an affine function of the observed outputs; thus, the affine controller can be run online at very high speeds, if needed. In contrast, MPC methods require solving an optimization problem at each time step, *i.e.*, online.

The affine controller design problem is not convex in its natural form, but can be transformed to an equivalent convex optimization problem by a nonlinear change of variables, which allows us to efficiently solve the problem. Our method is related to the classical Q -design procedure, or Youla parametrization [76], [77] for time-invariant, infinite-horizon linear controller design [1], [32], [33], [45], [47], [49], [73]. The book [20] and survey paper [21] describes this method, and the use of convex optimization to design continuous-time, time-invariant controllers, in detail; the Notes and References trace the ideas back into the 1960s. The books [30], [34] use these methods to minimize the worst-case output, with unknown but bounded input, which can be cast (after Q -parametrization) as an ℓ_1 norm minimization problem, and then solved using linear programming. The method is also related to the more recent purified output control method [7], [8]. While the method described in this paper is certainly related to the results given in these books and papers, the specifics are somewhat different.

The design of affine controllers in the context of discrete-time systems has received recent attention. It has been addressed in the context of robust optimization and robust controller design in [7], [8], [10], in the context of receding horizon control in [58], [59], for single-asset multistage investments with affine recourse and transaction costs in [67], for portfolio allocation problems with affine recourse in [24], for improved linear decision rules for stochastic programming in [27], and in the context of decentralized control over Banach spaces in [60].

We now mention what we believe is the closest work to the methods described in this paper, in the stochastic noise model setting. In the unpublished report [46], Hansson *et al.* describe a Q -parametrization method for computing an affine controller, for a *static* (*i.e.*, nondynamic) problem that satisfies some

second moment stochastic constraints, which can be expressed exactly (in the transformed variables) using second-order cone constraints. In a series of papers [68]–[71], van Hessem and Bosgra describe Q -parametrization methods for stochastic affine recourse MPC methods (as in our Section VIII). They approximate stochastic constraints, including chance constraints, as second moment constraints, which then translate into second-order cone constraints. Sung [64] applies their affine recourse MPC method to constrained portfolio optimization. Work similar in flavor to van Hessem's and Bosgra's was published by Wang *et al.* [74] and Goulart and Kerrigan [42] in which Q -parametrization is used to design affine recourse MPC methods for the special case of a convex quadratic objective function.

We also mention the closest work to the methods described in this paper, in the worst-case noise model setting. The idea that a Q -parametrization-like method can be used for minimax controller design in MPC was introduced in [51], [52]. It was then analyzed and extended in [41], [43], where Goulart *et al.* describe and analyze a Q -parametrization method for computing affine state feedback control policies in the context of linear discrete-time systems, which are subject to unknown but bounded state disturbances and mixed constraints on the state and input. They also show that such a formulation results in a stabilizing control law.

As described in the previous paragraphs, many of the ideas and techniques described in this paper are known, in general or other settings, such as infinite horizon, continuous-time time-invariant systems. Our contribution is to work out these ideas in detail, in a particular setting: discrete-time, finite horizon, time-varying system, with general convex objective and constraint functions, and arbitrary disturbance distribution. As far as we know, the case presented here has not appeared in the literature.

Although not directly related to our topic, we also mention two areas of research that share some similarities with our approach. Our method uses a (matrix) bilinear change of variables to convexify a non-convex controller design problem; this idea also came up in [39], [40], where a different change of variables is used to solve problems in the synthesis of state feedback gains, with LMI-based performance metric; see [22, §7.2.1]. The other area we mention is control-Lyapunov functions [36], [37], [56], [61], [62] in which a suboptimal control law can be found by convex optimization (in this case, over a set of approximate value functions).

Finally, we mention that there is a large literature on multi-stage stochastic linear programming, which can be used to solve (exactly or approximately) some versions of our problem (with piecewise linear objectives and polyhedral constraints); see [16], [17], [26], [31], [38], [44], [48], [65], [72]. The proposed methods range from decomposition and partitioning methods to sampling-based approximation algorithms, and are usually limited to short horizons.

In Sections II and III, we describe our model and the controller design problem in detail. We derive the formulas that relate the disturbances to the state and inputs in Section IV, and in Section V, we describe our convex formulation of the controller design problem. In the next two sections we illustrate the method with applications: a supply chain optimization problem in Section VI, and a dynamic portfolio optimization problem in Section VII. In Section VIII, we explain how the affine con-

troller synthesis method can be combined with a model predictive control technique.

II. SYSTEM MODEL

We consider a discrete-time linear time-varying system (or plant), over the time interval $t = 0, 1, \dots, T$, with dynamics

$$x(t+1) = A(t)x(t) + B(t)u(t) + w(t), \quad t = 0, \dots, T-1 \quad (1)$$

where $x(t) \in \mathbf{R}^n$ is the system state, $w(t) \in \mathbf{R}^n$ is the process noise or exogenous input, and $u(t) \in \mathbf{R}^m$ is the control input, all at time step t . The sensor or measured output is given by

$$y(t) = C(t)x(t) + v(t), \quad t = 0, \dots, T-1 \quad (2)$$

where $y(t) \in \mathbf{R}^p$, and $v(t)$ is the sensor noise at time period t . We will also be interested in the special case of *perfect state measurements*, in which $C(t) = I$ and $v(t) = 0$, i.e., $y(t) = x(t)$, for $t = 0, \dots, T-1$.

With slight abuse of notation, we will use x, w, u, y , and v to denote the associated trajectories

$$\begin{aligned} x &= (x(0), \dots, x(T)) \in \mathbf{R}^{n(T+1)} \\ u &= (u(0), \dots, u(T-1)) \in \mathbf{R}^{mT} \\ w &= (w(0), \dots, w(T-1)) \in \mathbf{R}^{nT} \\ y &= (y(0), \dots, y(T-1)) \in \mathbf{R}^{pT} \\ v &= (v(0), \dots, v(T-1)) \in \mathbf{R}^{pT}. \end{aligned}$$

We define $A_t^t = A(t-1)A(t-2) \cdots A(t)$, with $A_t^t = I$, and

$$C = [\text{diag}(C(0), \dots, C(T-1)) \quad 0] \in \mathbf{R}^{pT \times n(T+1)}$$

where $\text{diag}(\cdot)$ forms a block diagonal matrix from its arguments. The system (1) and (2) can then be written as

$$x = Gw + Hu + x_0, \quad y = Cx + v \quad (3)$$

where

$$x_0 = (x(0), A_0^1 x(0), \dots, A_0^T x(0)),$$

$$G = \begin{bmatrix} 0 & & & & & & \\ A_1^1 & 0 & & & & & \\ A_1^2 & A_2^2 & 0 & & & & \\ \vdots & & & \ddots & & & \\ \vdots & & & & & & 0 \\ A_1^T & A_2^T & \cdots & \cdots & A_T^T & & \end{bmatrix},$$

and

$$H = \begin{bmatrix} 0 & & & & & & \\ A_1^1 B(0) & 0 & & & & & \\ A_1^2 B(0) & A_2^2 B(1) & 0 & & & & \\ \vdots & & & \ddots & & & \\ \vdots & & & & & & 0 \\ A_1^T B(0) & A_2^T B(1) & \cdots & \cdots & A_T^T B(T-1) & & \end{bmatrix}.$$

(Entries not shown in these matrices are zero.)

We say that a matrix is (k, l) *block lower triangular* if it can be written as a block matrix with each block $k \times l$, and all blocks above the diagonal vanish. If in addition the diagonal blocks

vanish, we say the matrix is (k, l) *block strictly lower triangular*. (For $k = l = 1$, these correspond to the usual lower triangular and strictly lower triangular matrices.) We note for future use that G is (n, n) block strictly lower triangular, and H is (n, m) block strictly lower triangular.

III. CONTROLLER DESIGN

A. Causal and Affine Controllers

We will consider causal output feedback controllers, for which $u(t)$ has the form

$$u(t) = \varphi_t(y(0), \dots, y(t)), \quad t = 0, \dots, T-1.$$

The family of functions $\varphi_t : \mathbf{R}^{p(t+1)} \rightarrow \mathbf{R}^m$, for $t = 0, \dots, T-1$, is called the *control policy*. We will focus on the special case of affine causal output feedback, in which the functions φ_t are affine, i.e., have the form

$$u(t) = \varphi_t(y(0), \dots, y(t)) \quad (4)$$

$$= u_0(t) + \sum_{\tau=0}^t F(t, \tau)y(\tau), \quad t = 0, \dots, T-1. \quad (5)$$

We refer to $u_0(t) \in \mathbf{R}^m$ as the constant or open-loop part of the control, and we refer to $F(t, \tau) \in \mathbf{R}^{m \times p}$ as the (t, τ) feedback control gain. We interpret $u_0(t)$ as the input we apply if all previous measurements were zero, and we interpret $F(t, \tau)$ as the control or recourse gain from measurement $y(\tau)$ to input $u(t)$. If $F(t, \tau)$ is small, for example, it means that our choice of $u(t)$ is not heavily influenced by the measurement $y(\tau)$.

We can also express the control law in 'shifted form'. Let $\bar{y}(t)$ be some (arbitrary) nominal value of the output signal. We can express the control law (4) as

$$u(t) = \bar{u}(t) + \sum_{\tau=0}^t F(t, \tau)(y(\tau) - \bar{y}(\tau))$$

for $t = 0, \dots, T-1$, where

$$\bar{u}(t) = u_0(t) + \sum_{\tau=0}^t F(t, \tau)\bar{y}(\tau), \quad t = 0, \dots, T-1.$$

In this form we interpret $y(t) - \bar{y}(t)$ as the discrepancy between the measurement at time t and the nominal output value at time t . We interpret $\bar{u}(t)$ as the input we apply at time t if all previous discrepancies were zero, i.e., all previous measurements were equal to their nominal values. We interpret $F(t, \tau)$ as the gain from past discrepancies to the current input.

In the perfect state measurement case, the controller reduces to an affine causal state feedback controller, for which $u(t)$ is

$$u(t) = u_0(t) + \sum_{\tau=0}^t K(t, \tau)x(\tau), \quad t = 0, \dots, T-1.$$

Refer to $K(t, \tau) \in \mathbf{R}^{m \times n}$ as (t, τ) state feedback control gain.

Once we fix the control policy, the input trajectory u and the state trajectory x are functions of w and v , the process and measurement noises. We will consider two generic types of problems, which differ in how the unknown process and noise disturbances are modeled.

B. Stochastic Control

In the stochastic model, the process and sensor noises are random variables, with some known joint distribution. As a result, the input and state trajectories are also random variables. Let $\phi : \mathbf{R}^{n(T+1)} \times \mathbf{R}^{mT} \rightarrow \mathbf{R}$ be a convex objective function of the state and input trajectories. We judge our control performance by the expected value of this objective function

$$\mathbf{E}\phi(x, u)$$

where the expectation is over w and v . We treat constraints in a similar way. Let $\psi_i : \mathbf{R}^{n(T+1)} \times \mathbf{R}^{mT} \rightarrow \mathbf{R}$, $i = 1, \dots, M$, be a set of convex constraint functions. Our control design constraints are

$$\mathbf{E}\psi_i(x, u) \leq 0, \quad i = 1, \dots, M.$$

Such constraints, which require the expected value of a function to be less than zero (say) are called *stochastic constraints*. But the same form can be used to enforce an *almost-sure constraint*, such as $(x, u) \in \mathcal{C}$ almost surely. Here we simply define $\psi(x, u) = 0$ for $(x, u) \in \mathcal{C}$, and $\psi(x, u) = \infty$ for $(x, u) \notin \mathcal{C}$; the stochastic constraint $\mathbf{E}\psi(x, u) \leq 0$ is then equivalent to $(x, u) \in \mathcal{C}$ almost surely.

The stochastic controller design problem are expressed as

$$\begin{aligned} & \text{minimize} && \mathbf{E}\phi(x, u) \\ & \text{subject to} && \mathbf{E}\psi_i(x, u) \leq 0, \quad i = 1, \dots, M. \end{aligned} \quad (6)$$

In the general case of causal output feedback control, the optimization variables are the control policy functions φ_t , so the controller design problem is infinite dimensional. In the special case of affine causal feedback controllers, the optimization variables are $F(t, \tau)$ and u_0 . In this case the problem is finite dimensional. Here $u_0 = (u_0(0), \dots, u_0(T-1)) \in \mathbf{R}^{mT}$.

C. Worst-Case Robust Control

We can also model w and v using an unknown-but-bounded model. In this model, we are given a set \mathcal{Z} of possible values of (w, v) , and we judge the objective by the worst (*i.e.*, largest) value of a convex function over the set of possible disturbances

$$\sup_{(w,v) \in \mathcal{Z}} \phi(x, u).$$

The constraints are handled the same way: we require

$$\sup_{(w,v) \in \mathcal{Z}} \psi_i(x, u) \leq 0, \quad i = 1, \dots, M.$$

The worst-case robust controller design problem is thus

$$\begin{aligned} & \text{minimize} && \sup_{(w,v) \in \mathcal{Z}} \phi(x, u) \\ & \text{subject to} && \sup_{(w,v) \in \mathcal{Z}} \psi_i(x, u) \leq 0, \quad i = 1, \dots, M. \end{aligned} \quad (7)$$

As in the stochastic controller design problem, the variables are the policy functions φ_t (in the general case), or the matrices $F(t, \tau)$ and the vector u_0 (in the affine controller case).

IV. CLOSED-LOOP SYSTEM

We now derive the explicit dependence of x and u on w and v , when the controller is affine. We define feedback matrix as

$$F = \begin{bmatrix} F(0,0) & & & & \\ F(1,0) & F(1,1) & & & \\ \vdots & & \ddots & & \\ F(T-1,0) & F(T-1,1) & \cdots & F(T-1,T-1) & \end{bmatrix}$$

which is (m, p) block lower triangular. We then have

$$u = Fy + u_0. \quad (8)$$

From (3) and (8) we can solve for x and u in terms of w and v , to get

$$\begin{bmatrix} x \\ u \end{bmatrix} = P \begin{bmatrix} w \\ v \end{bmatrix} + \begin{bmatrix} \tilde{x} \\ \tilde{u} \end{bmatrix} \quad (9)$$

where

$$\begin{aligned} P &= \begin{bmatrix} P_{xw} & P_{xv} \\ P_{uw} & P_{uv} \end{bmatrix}, \\ P_{xw} &= G + HF(I - CHF)^{-1}CG \\ P_{xv} &= HF(I - CHF)^{-1} \\ P_{uw} &= F(I - CHF)^{-1}CG \\ P_{uv} &= F(I - CHF)^{-1}, \\ & \text{and} \\ \tilde{x} &= x_0 + Hu_0 + HF(I - CHF)^{-1}(Cx_0 + CHu_0) \\ \tilde{u} &= F(I - CHF)^{-1}(Cx_0 + CHu_0) + u_0. \end{aligned}$$

The matrix CHF is (p, p) block strictly lower triangular (therefore also strictly lower triangular), so $I - CHF$ is invertible.

We refer to P as the *closed-loop matrix*, \tilde{x} as the *closed-loop state trajectory*, and \tilde{u} as the *closed-loop control trajectory*. Evidently \tilde{x} and \tilde{u} are the (closed-loop) state and input trajectories when the process and noise disturbances are zero. The matrices P_{xw} and P_{xv} tell us how sensitive the state trajectory is to the process and sensor noises, respectively. The matrices P_{uw} and P_{uv} tell us how sensitive the input trajectory is to the process and sensor noises.

Using (9), we can express the objective and constraints in terms of the closed-loop quantities P , \tilde{x} , and \tilde{u} . In the stochastic model, the objective is

$$\mathbf{E}\phi(x, u) = \mathbf{E}\phi(\tilde{x} + P_{xw}w + P_{xv}v, \tilde{u} + P_{uw}w + P_{uv}v)$$

which is a convex function of P , \tilde{x} , and \tilde{u} . To see this, we note that for any particular realizations of w and v , $\phi(\tilde{x} + P_{xw}w + P_{xv}v, \tilde{u} + P_{uw}w + P_{uv}v)$ is a convex function of the closed-loop quantities; the expectation of a family of convex function is convex [23, §3.2.1].

The constraint functions can be written as

$$\mathbf{E}\psi_i(x, u) = \mathbf{E}\psi_i(\tilde{x} + P_{xw}w + P_{xv}v, \tilde{u} + P_{uw}w + P_{uv}v)$$

which, like the objective function, are convex functions of P , \tilde{x} , and \tilde{u} . It follows that the stochastic controller design problem (6) is convex in the closed-loop quantities P , \tilde{x} , and \tilde{u} .

In the worst-case robust model, the objective function is

$$\sup_{(w,v) \in \mathcal{Z}} \phi(\tilde{x} + P_{xw}w + P_{xv}v, \tilde{u} + P_{uw}w + P_{uv}v)$$

and the constraint functions are

$$\sup_{(w,v) \in \mathcal{Z}} \psi_i(\tilde{x} + P_{xw}w + P_{xv}v, \tilde{u} + P_{uw}w + P_{uv}v).$$

These are convex functions of P , \tilde{x} , and \tilde{u} , since the supremum over a family of convex functions is convex [23, §3.2.3]. Thus, the worst-case robust controller design problem (7) is convex in the closed-loop quantities P , \tilde{x} , and \tilde{u} .

A. Perfect State Measurement Case

In the perfect state measurement case we have $u = u_0 + Kx$, where K is an (m, n) block lower triangular matrix

$$K = \begin{bmatrix} K(0,0) & & & & 0 \\ K(1,0) & K(1,1) & & & 0 \\ \vdots & & \ddots & & 0 \\ K(T-1,0) & K(T-1,1) & \cdots & K(T-1,T-1) & 0 \end{bmatrix}.$$

We then have

$$\begin{bmatrix} x \\ u \end{bmatrix} = Pw + \begin{bmatrix} \tilde{x} \\ \tilde{u} \end{bmatrix} \quad (10)$$

where

$$P = \begin{bmatrix} P_{xw} \\ P_{uw} \end{bmatrix},$$

$$P_{xw} = G + HK(I - HK)^{-1}G$$

$$P_{uw} = K(I - HK)^{-1}G,$$

and

$$\tilde{x} = x_0 + Hu_0 + HK(I - HK)^{-1}(x_0 + Hu_0)$$

$$\tilde{u} = K(I - HK)^{-1}(x_0 + Hu_0) + u_0.$$

V. CONVEX FORMULATION

Problems (6) and (7) are in general not convex in our design variables F and u_0 . By a suitable change of variables, however, these problems can be cast as convex optimization problems, and therefore solved efficiently.

We define

$$Q = F(I - CHF)^{-1} = P_{uw} \quad (11)$$

the closed-loop matrix from v to u . The matrix Q is (m, p) block lower triangular. Conversely, suppose that Q is any (m, p) block lower triangular matrix in $\mathbf{R}^{mT} \times \mathbf{R}^{pT}$. We can then define

$$F = (I + QCH)^{-1}Q. \quad (12)$$

First we note that QCH is (m, p) block strictly lower triangular, so $I + QCH$ is invertible and (m, m) block lower triangular; its inverse is therefore also (m, m) block lower triangular, and thus F in (12) is (m, p) block lower triangular. Thus, (11) and (12) define a bijection between (m, p) block lower triangular F and Q . In particular, we can use the variable Q , restricted to be (m, p) block lower triangular, in place of F .

We also define

$$r = (I + QCH)u_0 \quad (13)$$

so that

$$u_0 = (I + QCH)^{-1}r = (I + FCH)r. \quad (14)$$

We will use the variable r instead of u_0 .

We can express the closed-loop matrices in terms of Q , and the closed-loop responses in terms of Q and r as

$$P = \begin{bmatrix} P_{xw} & P_{xv} \\ P_{uw} & P_{uv} \end{bmatrix} = \begin{bmatrix} (I + HQC)G & HQ \\ QCG & Q \end{bmatrix},$$

$$\tilde{x} = (I + HQC)x_0 + Hr,$$

$$\tilde{u} = QCx_0 + r.$$

Now we make the critical observation that the closed-loop quantities P , \tilde{x} , and \tilde{u} are *affine functions* of Q and r . It follows that the problems (6) and (7) are convex in the variable Q (restricted to be (m, p) block lower triangular) and r .

The stochastic problem (6) can be cast as

$$\begin{aligned} &\text{minimize} && \mathbf{E}\phi(\tilde{x} + P_{xw}w + P_{xv}v, \\ &&& \tilde{u} + P_{uw}w + P_{uv}v) \\ &\text{subject to} && \mathbf{E}\psi_i(\tilde{x} + P_{xw}w + P_{xv}v, \\ &&& \tilde{u} + P_{uw}w + P_{uv}v) \leq 0, \\ &&& i = 1, \dots, M, \\ &&& Q(m, p) \text{ block lower triangular.} \end{aligned} \quad (15)$$

Similarly the worst-case problem (7) can be cast as

$$\begin{aligned} &\text{minimize} && \sup_{(w,v) \in \mathcal{Z}} \phi(\tilde{x} + P_{xw}w + P_{xv}v, \\ &&& \tilde{u} + P_{uw}w + P_{uv}v) \\ &\text{subject to} && \sup_{(w,v) \in \mathcal{Z}} \psi_i(\tilde{x} + P_{xw}w + P_{xv}v, \\ &&& \tilde{u} + P_{uw}w + P_{uv}v) \leq 0, \\ &&& i = 1, \dots, M, \\ &&& Q(m, p) \text{ block lower triangular.} \end{aligned} \quad (16)$$

The variables in (15) and (16) are Q and r ; the expressions \tilde{x} , \tilde{u} , P_{xw} , P_{xv} , P_{uw} , P_{uv} are affine functions of these variables. Once we solve the problem (15) or (16), we can recover our original design variables F and u_0 using (12) and (14).

A. Perfect State Measurement Case

In the perfect state measurement case we define

$$Q = K(I - HK)^{-1}$$

where $Q \in \mathbf{R}^{mT \times n(T+1)}$ is (m, n) block lower triangular. We can then define

$$F = (I + QH)^{-1}, \quad r = (I + QH)u_0 \quad (17)$$

and recover u_0 as

$$u_0 = (I + QH)^{-1}r = (I + KH)r. \quad (18)$$

We can express the closed-loop matrices in terms of Q , and the closed-loop responses in terms of Q and r as

$$P = \begin{bmatrix} P_{xw} \\ P_{uw} \end{bmatrix} = \begin{bmatrix} (I + HQ)G \\ QG \end{bmatrix}$$

$$\tilde{x} = (I + HQ)x_0 + Hr,$$

$$\tilde{u} = Qx_0 + r.$$

The stochastic problem (6) can be cast as

$$\begin{aligned} & \text{minimize} && \mathbf{E}\phi(\tilde{x} + P_{xw}w, \\ & && \tilde{u} + P_{uw}w) \\ \text{subject to} && \mathbf{E}\psi_i(\tilde{x} + P_{xw}w, \\ && \tilde{u} + P_{uw}w) \leq 0, \\ && i = 1, \dots, M, \\ && Q(m, p) \text{ block lower triangular.} \end{aligned} \quad (19)$$

Similarly the worst-case problem (7) can be cast as

$$\begin{aligned} & \text{minimize} && \sup_{w \in \mathcal{Z}} \phi(\tilde{x} + P_{xw}w, \\ & && \tilde{u} + P_{uw}w) \\ \text{subject to} && \sup_{w \in \mathcal{Z}} \psi_i(\tilde{x} + P_{xw}w, \\ && \tilde{u} + P_{uw}w) \leq 0, \\ && i = 1, \dots, M, \\ && Q(m, p) \text{ block lower triangular.} \end{aligned} \quad (20)$$

The variables in (19) and (20) are Q and r ; the expressions \tilde{x} , \tilde{u} , P_{xw} , P_{uw} , are affine functions of these variables. Once we solve the problem (19) or (20), we can recover our original design variables K and u_0 using (17) and (18).

B. Solving the Convex Controller Design Problem

The change of variables described above yields a convex stochastic program [in the case of (6)], or a robust convex optimization problem [in the case of (7)]. In some cases, these problems can be solved by formulating them as standard convex optimization problems, with no expectation or supremum. Simple examples include the case where ϕ and ψ_i are (convex) quadratic functions, and (w, v) is Gaussian (in the stochastic case), or known to lie in an ellipsoid (in the worst-case framework).

In other cases we can solve (6) and (7) approximately. In the stochastic framework we can resort to sampling, *i.e.*, replacing expectations by empirical means with values chosen randomly from the given distribution. For more on stochastic programming and stochastic control see [14], [18], [75]. In the worst-case framework, we resort to robust convex optimization techniques, both exact and approximate (see, *e.g.*, [7]–[9], [11], [25], [54], [78]).

C. Block Banded Q Design

We say that a matrix is (k, l) block banded with bandwidth $2b + 1$ if it can be written as a block matrix with each block $k \times l$, and all blocks above the b th upper diagonal and below the b th lower diagonal vanish. If all blocks above the main diagonal and below the b th lower diagonal vanish, the matrix is (k, l) block lower triangular and block banded with bandwidth b . The

bandwidth b can vary between 1 (*i.e.*, Q is block diagonal) and T (*i.e.*, Q is a full block lower triangular matrix).

We add the constraint that Q is (m, p) block banded (Q is then block banded and block lower triangular). This does not mean the controller has bounded memory; it means the closed-loop response from v to u does. It is, of course, a convex constraint on Q , and so can be handled as above.

Since we are restricting the possible values of Q , imposing this condition evidently increases the objective, *i.e.*, results in poorer performance. Why would we do this? The main reason is it reduces the size of the convex problem that we need to solve, sometimes dramatically. For example, if we only constrain Q to be (m, p) lower triangular, the number of variables in (15) and (16) is $(T^2 + T)mp/2$, which grows quadratically with T . However, if we also constrain Q to be (m, p) banded, with bandwidth b , the number of variables in (15) and (16) is $(T - b + 1)bmp/2$, which grows linearly with T .

VI. SUPPLY CHAIN OPTIMIZATION

We consider a single (divisible) commodity linear supply chain problem, where, in each period, we can ship the commodity between a store and a factory to meet a (random) demand. Our goal is to minimize the total cost incurred, taking into consideration shipping cost, storage cost, and backlog cost.

We let $x(t) \in \mathbf{R}$ be the amount of the commodity available at the store at time t , for $t = 0, \dots, T$, with $x(t) < 0$ meaning a backlog of $|x(t)|$ units of the commodity. The demand for the commodity at time t is denoted $d(t)$. The amount of the commodity shipped to the store at time t is denoted $u(t)$, with $u(t) < 0$ meaning the amount $|u(t)|$ is shipped back to the factory from the store. The store inventory updates as

$$x(t + 1) = x(t) + u(t) + w(t) \quad (21)$$

where $w(t) = -d(t)$.

The shipping cost is proportional to the amount shipped, *i.e.*, $s|u(t)|$, where s is the shipping rate. The storage cost is proportional to the inventory $x(t)$, when it is positive, *i.e.*, $hx(t)$. When the inventory is negative, *i.e.*, represents a backlog, the cost is $-bx(t)$. (This might represent the cost of processing the backlog, customer ill-will, etc.) The total cost incurred in time period t , for $t = 0, \dots, T - 1$, is

$$\phi(x(t), u(t)) = \max(hx(t), -bx(t)) + s|u(t)|.$$

At time $t = T$, the cost incurred is

$$\phi_{\text{final}}(x(T)) = \max(gx(T), -bx(T))$$

where g is the salvage cost.

Our objective is to minimize the expected total cost, *i.e.*

$$\text{minimize } \mathbf{E} \left(\sum_{t=0}^{T-1} \phi(x(t), u(t)) + \phi_{\text{final}}(x(T)) \right) \quad (22)$$

where expectation is taken over the demand d .

When the demands $d(t)$, $t = 0, \dots, T - 1$, are independent, the supply chain optimization problem (22) has an analytical solution, and the optimal controller is an (s, S) policy (see [14, §4.2] for more details.) When the demands $d(t)$ are correlated across time, however, there is no general analytical solution.

TABLE I
SIMULATION RESULTS FOR SUPPLY CHAIN EXAMPLE

Control law	Cost mean \pm std. dev.
Optimal affine	185.6 \pm 63.2
Naïve greedy	222.4 \pm 71.3
Greedy	208.7 \pm 68.7
CE-MPC	183.3 \pm 62.6
Prescient	(165.9 \pm 53.0)

We can find the optimal affine controller for (22) using several methods, depending on the particular distribution of d . If d has a discrete distribution, and can take on only finitely many values (demand scenarios), the affine controller design problem can be reduced to an LP and solved exactly. However, if d has a continuous distribution, the affine controller design problem will have to be approximated by sampling from the distribution, or by other stochastic optimization methods. If we replace the expectation in (22) by the empirical mean over M sample demand trajectories $d^{(j)}$, $j = 1, \dots, M$, the problem becomes

$$\text{minimize } \left(\frac{1}{M} \right) \sum_{j=1}^M \left(\sum_{t=0}^{T-1} \phi(x^{(j)}(t), u^{(j)}(t)) + \phi_{\text{final}}(x^{(j)}(T)) \right) \quad (23)$$

where the variables are still Q and r , and $x^{(j)}(t)$, $u^{(j)}(t)$ correspond to the j th state and control input trajectories at time t (corresponding to the j th sample $d^{(j)}$).

We can also find a lower bound on the optimal value of the stochastic inventory control problem by assuming the controller is prescient, *i.e.*, that the entire demand trajectory d is known at time 0. The inventory updates (21) are then deterministic and the optimal control trajectory u , given d , can be found by solving the (deterministic) optimization problem

$$\begin{aligned} & \text{minimize } \left(\sum_{t=0}^{T-1} \phi(x(t), u(t)) + \phi_{\text{final}}(x(T)) \right) \\ & \text{subject to } x(t+1) = x(t) + u(t) - d(t), \\ & \quad t = 0, \dots, T-1. \end{aligned} \quad (24)$$

The expected value of the optimal value of this problem, over the distribution of d , gives a lower bound on the optimal value of the inventory control problem, for any causal controller. We can (accurately) estimate this by generating demand samples and taking the average of the optimal values.

A. Example

We consider the case where d has a lognormal distribution, *i.e.*, where $\log(d) \sim \mathcal{N}(\mu, \Sigma)$. Here $\log(\cdot)$ is applied entrywise, $\mu \in \mathbf{R}^T$, and $\Sigma \in \mathbf{R}^{T \times T}$. The conditional distribution of $d(t)$ given $d(0), \dots, d(t-1)$ is also log normal, *i.e.*,

$$\log(d(t)|d(0), \dots, d(t-1)) \sim \mathcal{N}(\nu(t), \Lambda(t))$$

for $t = 1, \dots, T-1$. Here

$$\begin{aligned} \nu(t) &= \mu_t + \Sigma_{t,0:t-1} \Sigma_{0:t-1,0:t-1}^{-1} \delta(t) \\ \delta(t) &= \log((d(0), \dots, d(t-1))) - \mu_{0:t-1}, \\ \Lambda(t) &= \Sigma_{t,t} - \Sigma_{t,0:t-1} \Sigma_{0:t-1,0:t-1}^{-1} \Sigma_{t,0:t-1}^T \end{aligned}$$

where μ_t and $\Sigma_{t,t}$ are, respectively, the mean and covariance of $\log(d(t))$, $\mu_{0:t-1}$ and $\Sigma_{0:t-1,0:t-1}$ are, respectively, the mean and covariance of $\log((d(0), \dots, d(t-1)))$, and $\Sigma_{t,0:t-1}$ is

the covariance of $d(t)$ and $(d(0), \dots, d(t-1))$. Note that all the mentioned vectors are subvectors of μ and the mentioned matrices are submatrices of Σ .

We will compare the optimal affine controller obtained from solving (22) to several rival approaches. We first consider the naïve greedy controller which assumes that the demand at time t will be its mean

$$\bar{d}(t) = \mathbf{E}(d(t))$$

and chooses $u(t)$ to minimize the corresponding stage cost at $t+1$. The naïve greedy control law is therefore

$$u(t) = \bar{d}(t) - x(t), \quad t = 0, \dots, T-1.$$

Note that the naïve greedy control law is affine in x and therefore will perform worse than the (optimal) affine controller obtained from solving (22).

The second controller we consider is the greedy controller that picks $u(t)$ so as to minimize the expected stage cost at $t+1$. Let $\hat{d}(t)$ be

$$\hat{d}(t) = \mathbf{E}(d(t)|d(0), \dots, d(t-1))$$

for $t = 1, \dots, T-1$. We take $\hat{d}(0) = \mathbf{E}(d(0)) = \mu_0$. The greedy control law is therefore

$$u(t) = \hat{d}(t) - x(t), \quad t = 0, \dots, T-1.$$

Note that $\hat{d}(t)$ is *not* affine in $d(0), \dots, d(t-1)$ and therefore *not* affine in $x(t)$. The greedy control law is therefore *not* affine in x .

The last rival approach that we consider is certainty-equivalent model predictive control (CE-MPC), explained in detail in Section VIII.

We take $T = 10$, $s = 8$, $h = 1$, $b = 4$, $g = 4$, and we generate μ and Σ randomly: μ_i are independent identically distributed (IID) on $[-1.5, 1.5]$, and $\Sigma = KK^T$ where $K \in \mathbf{R}^{T \times T}$ and K_{ij} IID $\mathcal{N}(0, 0.25)$.

We first generate 1000 samples from the distribution of d and solve (23). We then generate another 5000 samples from the distribution of d and test the performance of the affine controller versus the performances of the three rival methods. The results are shown in Table I, along with the performance of the (non-causal) prescient controller, which provides a lower on achievable performance. The optimal affine controller beats the two greedy controllers by a substantial margin, and has performance comparable to that of CE-MPC, and only around 12% higher than the lower bound given by the (non-causal) prescient controller.

Fig. 1 shows the histograms of the costs achieved by the optimal affine controller, the naïve greedy controller, the greedy controller, CE-MPC, and, finally, the prescient controller.

VII. DYNAMIC PORTFOLIO OPTIMIZATION

We consider a dynamic portfolio optimization problem, in which we buy and sell assets each period, in order to generate an income stream, while at the same time keeping our portfolio balanced, *i.e.*, not too far from some target portfolio.

We let $x_i(t)$ be our holding (in dollars) in asset i in period t , with $x_i(t) < 0$ meaning a short position, for $i = 1, \dots, n$, $t = 0, \dots, T$. We let $x^{\text{tar}} \in \mathbf{R}^n$ be our target or desired portfolio;




Fig. 1. Histograms of the total cost incurred using (a) the naïve greedy controller, (b) the greedy controller, (c) CE-MPC, (d) the optimal affine controller, and (e) the (non-causal) prescient control law.

one of our goals is to maintain $x(t) \approx x^{\text{tar}}$, despite fluctuations in returns. Our holdings update as

$$x_i(t+1) = (1 + r_i(t))(x_i(t) + u_i(t)) \quad (25)$$

where $r_i(t)$ is the return on asset i in period t , and $u_i(t)$ is amount of asset i bought at the beginning of the period (or sold, if $u_i(t) < 0$). The return vectors $r(t)$ are IID, from some known distribution, with $r_i(t) > -1$ almost surely, and mean $\bar{r} = \mathbf{E}r(t) \geq 0$. We will refer to (25) as the nonlinear update equations.

We will make some approximations to put the update (25) into our linear dynamical system form. First we write (25) as

$$x_i(t+1) = (1 + \bar{r}_i)(x_i(t) + u_i(t)) + (r_i(t) - \bar{r}_i)(x_i^{\text{tar}} + (x_i(t) - x_i^{\text{tar}}) + u_i(t)).$$

We now assume that $u_i(t)$ and $x_i(t) - x_i^{\text{tar}}$ are small compared to $x_i(t)$. Dropping these terms we have the approximate or linearized update equations

$$x_i(t+1) = (1 + \bar{r}_i)(x_i(t) + u_i(t)) + w_i(t) \quad (26)$$

$$w_i(t) = (r_i(t) - \bar{r}_i)x_i^{\text{tar}}. \quad (27)$$

Here $w(t)$ can be thought of as the absolute return vector (since its entries are in dollars, not percentages); these are IID, from a known distribution. We will refer to (26) and (27) as the linearized update equations in the sequel.

In addition to requiring $x(t)$ to be close to x^{tar} in each time period t , we also require that the budget be conserved at each step, in expectation, *i.e.*, that $\mathbf{E}\mathbf{1}^T x(t) = \mathbf{1}^T x^{\text{tar}}$ for $t = 0, \dots, T$.

At the beginning of time period t , the net cash taken out of the portfolio (or put into it, if it is negative) is $-\mathbf{1}^T u(t)$. From this we subtract a fee for transaction costs, proportional to the absolute value of the amount of each asset bought or sold, so the total cash taken out of the portfolio (*i.e.*, the income) at time t is

$$-\mathbf{1}^T u(t) - \kappa \|u(t)\|_1$$

where $\kappa \in (0, 1)$ is the proportional transaction fee rate. We will measure the utility of this income with a utility function $\phi: \mathbf{R} \rightarrow \mathbf{R}$, which is increasing and concave; the total utility is

$$U(u) = \sum_{t=0}^{T-1} \phi(-\mathbf{1}^T u(t) - \kappa \|u(t)\|_1)$$

which is a concave function of u (using one of the composition rules in convex analysis; see [23, §3.2.4]).

Now we can formulate our problem as

$$\begin{aligned} & \text{maximize} && \mathbf{E}U(u) \\ & \text{subject to} && \mathbf{E}\|x(t) - x^{\text{tar}}\|_2^2 \leq n\rho^2, \quad t = 0, \dots, T, \\ & && \mathbf{E}\mathbf{1}^T x(t) = \mathbf{1}^T x^{\text{tar}}, \quad t = 0, \dots, T. \end{aligned} \quad (28)$$

Here $\rho > 0$ is a parameter that controls how far from the target portfolio we allow our portfolio to drift. This problem is a standard convex stochastic control problem, so we can determine the optimal affine causal feedback control law using the methods described above.

We can find the optimal affine controller for (28) using several methods. If $w(t)$ has a discrete distribution (*i.e.*, we model the returns as taking one of N values, with some given probabilities) then we can solve it using standard convex optimization techniques. In most cases, however, $r(t)$, and therefore also $w(t)$, will have a continuous distribution. Because of the ℓ_1 norm appearing in the objective function, it is very unlikely that the objective can be expressed analytically; thus, it will have to be approximated by, say, sampling from the distribution, or any other stochastic optimization method. The constraints, however, can be expressed exactly using second-order cone (or convex quadratic) constraints, since $\mathbf{E}\|x(t) - x^{\text{tar}}\|_2^2$ is a convex quadratic function of F and r .

If we are given the covariance Σ_r of $r(t)$, and if we replace the expectation in the objective of (28) by the empirical mean over M sample trajectories $r^{(j)}$, $j = 1, \dots, M$, the problem is

$$\begin{aligned} & \text{maximize} && \left(\frac{1}{M}\right) \sum_{j=1}^M U(u^{(j)}) \\ & \text{subject to} && \|Z_t P_x w S^{1/2}\|_F^2 + \|Z_t \tilde{x} - x^{\text{tar}}\|_2^2 \leq n\rho^2 \\ & && t = 0, \dots, T, \\ & && \mathbf{1}^T \tilde{x}(t) = \mathbf{1}^T x^{\text{tar}}, \quad t = 0, \dots, T \end{aligned} \quad (29)$$

where $x^{(j)}$ and $u^{(j)}$ are the j th state and control input trajectories, corresponding to the j th sample $r^{(j)}$, Z_t is a $n \times n(T+1)$ matrix whose $(t+1)$ th $n \times n$ block is the identity and

$$S = \text{diag}(\Sigma_w, \dots, \Sigma_w) \in \mathbf{R}^{n(T+1) \times n(T+1)}.$$

Here Σ_w is the covariance matrix of w and is given by

$$\Sigma_w = \text{diag}(x^{\text{tar}}) \Sigma_r \text{diag}(x^{\text{tar}}).$$

A. Example

Consider a particular problem instance with $n = 4$ assets, $T = 10$ periods, $\kappa = 0.1$ and $\rho = 0.1$. We choose the target portfolio to be $x^{\text{tar}} = (1/3, 0, 1/3, 1/3)$ and let $x(0) = x^{\text{tar}}$.

We model the returns $r(t)$ to be IID, with $(1 + r(t))$ lognormally distributed with mean and covariance

$$\begin{aligned} \bar{r} &= [0.0202 \quad 0.0489 \quad 0.0762 \quad 0.1012]^T, \\ \Sigma_r &= \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0.0045 & 0.0052 & 0.0034 \\ 0 & 0.0052 & 0.0158 & 0.0040 \\ 0 & 0.0034 & 0.0040 & 0.0312 \end{bmatrix}. \end{aligned}$$

(See Section A for details on the method used to generate \bar{r} and Σ_r .)

TABLE II
MEAN UTILITIES FOUND BY SIMULATION

Control law	Mean utility
Optimal affine, linearized dynamics	-3.8×10^{-3}
Optimal affine, exact dynamics	-3.8×10^{-3}
Greedy	-5.7×10^{-3}

We choose the utility function to be the negative quadratic downside risk

$$\phi(z) = -(r_{\text{des}} - z)_+^2$$

where $(z)_+ = \max(0, z)$. Here r_{des} is the desired portfolio return, which we take to be $r_{\text{des}} = 10\%$.

We compare our approach to a greedy nonlinear controller

$$u_i(t) = \frac{x_i^{\text{tar}}}{(1 + \bar{r}_i)} - x_i(t).$$

The greedy controller assumes that the returns will equal their means, and re-balances to the target portfolio in each time period. Note that under this scheme the budget conservation constraint is met.

We generate $M = 1000$ sample trajectories for r and solve the sampled problem (29) to obtain the affine controller parameters. We then test the resulting affine controller on 5000 sample trajectories of r , using the exact nonlinear dynamics (25), as well as the approximate linearized dynamics (26) and (27). We also test the greedy controller on the sample return trajectories.

The results are shown in Table II. There is virtually no difference between the results in the case of the affine controller used with approximate or exact dynamics, confirming that our linearization approximation is good. We can see that the affine controller substantially outperforms the greedy controller.

VIII. AFFINE RECOURSE MODEL PREDICTIVE CONTROL

In this section we describe an approximate solution for the stochastic control problem (6), that combines the affine control laws described above with the basic idea behind model predictive control.

A. Certainty-Equivalent MPC

The most common form of MPC is *certainty-equivalent MPC* (CE-MPC). In CE-MPC, the control $u(t)$ at time t is found as follows. First, we form estimates of the current and future process noises, $w(t), \dots, w(T-1)$ (which are not yet known), given the state information available at time t , *i.e.*, $x(0), \dots, x(t)$. (From these we can find $w(0), \dots, w(t-1)$.) One natural choice (but not the only one) is the conditional mean, *i.e.*, $\mathbf{E}(w(t)|w(0), \dots, w(t-1))$. We will denote these estimates as $\hat{w}(\tau|t)$, for $\tau = t, t+1, \dots, T-1$.

We now form and solve the (deterministic) convex optimization problem

$$\begin{aligned} & \text{minimize} && \phi(x, u) \\ & \text{subject to} && \psi_i(x, u) \leq 0, \quad i = 1, \dots, M, \\ & && \text{dynamics constraints} \end{aligned} \quad (30)$$

where we describe the dynamics constraint shortly, and we must carefully interpret the symbols x and u . The components $x(0), \dots, x(t)$ are the actual, measured current and past states; $x(t+1), \dots, x(T)$ are to be interpreted as variables in the opti-

TABLE III
SIMULATION RESULTS FOR SUPPLY CHAIN EXAMPLE

Control law	Cost mean \pm std. dev.
Optimal affine	185.6 ± 63.2
CE-MPC	183.3 ± 62.6
ARE-MPC	180.5 ± 61.9
Prescient	(165.9 ± 53.0)

mization problem. Likewise, $u(0), \dots, u(t-1)$ are interpreted as the actual previous choices of the input; $u(t), \dots, u(T-1)$ are interpreted as variables. Finally, we describe the dynamics constraints, which are

$$x(\tau+1) = A(\tau)x(\tau) + Bu(\tau) + \hat{w}(\tau|t)$$

for $\tau = t, \dots, T-1$. Once we solve this problem, we take $u(t)$ as our input. This whole process (re-estimating the current and future disturbances, forming and solving a new optimization problem) is repeated at time $t+1$.

CE-MPC involves two gross approximations, one of which makes the problem easier, and one harder. CE-MPC completely ignores all recourse in the future, since it replaces the future disturbances with deterministic estimates. In fact, we will have recourse, since we will obtain future states, and can act accordingly. On the other hand, we also ignore all statistical variation in the future, which Jensen's inequality tells us will increase the expected value of the objective (*i.e.*, when conditional means are used for the estimates). Despite these gross approximations, CE-MPC works surprisingly well in many applications.

B. Affine Recourse MPC

In this section, we explain how the synthesis of affine controllers, as described above, can be combined with MPC. We refer to this method as *affine-recourse MPC* (AR-MPC). The same idea has already been proposed by van Hessen and Bosgra [68]–[71] who consider only second moment constraints (by approximation, in some cases), and so end up with second-order cone programs to solve. The idea was also used by Goulart *et al.* [41]–[43] in the context of robust MPC.

In AR-MPC, our plan includes the idea of recourse in the future, but limits this recourse to be affine. The control $u(t)$ at time t is found as follows. First, we compute the conditional distribution of $w(t), \dots, w(T-1)$ given the state information available at time t , *i.e.*, $x(0), \dots, x(t)$. (This is the same as the distribution of $w(t), \dots, w(T-1)$ given $w(0), \dots, w(t-1)$.)

We assume that the input at times $t, \dots, T-1$ is affine in the current and future states, *i.e.*

$$u(t+s) = u_0^{(t)}(s) + \sum_{\tau=0}^s K^{(t)}(s, \tau)x(t+\tau) \quad (31)$$

for $s = 0, \dots, T-t-1$, and where $u_0^{(t)}(s) \in \mathbf{R}^n$ and $K^{(t)}(s, \tau) \in \mathbf{R}^{n \times n}$.

We now form the (stochastic) optimization problem

$$\begin{aligned} & \text{minimize} && \mathbf{E}\phi(x, u) \\ & \text{subject to} && \mathbf{E}\psi_i(x, u) \leq 0, \quad i = 1, \dots, M \end{aligned} \quad (32)$$

where the expectation is taken over the conditional distribution of $(w(t), \dots, w(T-1))$ given $w(0), \dots, w(t-1)$, and

we must carefully interpret the symbols x and u . The components $x(0), \dots, x(t)$ are the actual, measured current and past states; $u(0), \dots, u(t-1)$ are interpreted as the actual previous choices of the input. The components $x(t+1), \dots, x(T)$ and $u(t), \dots, u(T-1)$ are related by (31) and the optimization variables are $K^{(t)}(s, \tau)$ and $u_0^{(t)}(s)$. We solve (32) using the method presented in Section V and take

$$u(t) = u_0^{(t)}(0) + K^{(t)}(0, 0)x(t).$$

This whole process (re-estimating the conditional distribution of current and future disturbances, forming and solving a new stochastic optimization problem) is repeated at time $t+1$.

C. Example

We compare the performances of the optimal affine controller, CE-MPC and AR-MPC on the supply chain numerical example used in Section VI-A.

The conditional distribution of $d(t), \dots, d(T-1)$ given $d(0), \dots, d(t-1)$ is also lognormal, *i.e.*

$$\log(d(t), \dots, d(T-1)) | d(0), \dots, d(t-1) \sim \mathcal{N}(\nu(t), \Lambda(t)) \quad (33)$$

for $t = 1, \dots, T-1$. Here

$$\begin{aligned} \nu(t) &= \mu_{t:T-1} + \Sigma_{t:T-1, 0:t-1} \Sigma_{0:t-1, 0:t-1}^{-1}, \\ \delta(t) &= (\log((d(0), \dots, d(t-1)))) - \mu_{0:t-1}, \\ \Lambda(t) &= \Sigma_{t:T-1, t:T-1} - \Sigma_{t:T-1, 0:t-1} \Sigma_{0:t-1, 0:t-1}^{-1} \Sigma_{0:t-1, t:T-1} \end{aligned}$$

where $\mu_{t:T-1}$ and $\Sigma_{t:T-1, t:T-1}$ are, respectively, the mean and covariance of $\log((d(t), \dots, d(T-1)))$, $\mu_{0:t-1}$ and $\Sigma_{0:t-1, 0:t-1}$ are, respectively, the mean and covariance of $\log((d(0), \dots, d(t-1)))$, and $\Sigma_{t:T-1, 0:t-1}$ is the covariance of $(d(t), \dots, d(T-1))$ and $(d(0), \dots, d(t-1))$.

We test these 3 methods on 2000 sample paths from the distribution of d . The optimal affine controller is obtained by solving (23), using 500 other sample paths generated from the distribution of d . For each test sample path, we compute the CE-MPC control at each step by replacing future demand with the conditional mean, solving a deterministic optimization problem (30) and using only the first step of the resulting control plan as control. In the AR-MPC case, for each test sample path and at each time step t , we generate 500 samples of demand vectors from the conditional distribution (33) and solve the shortened-horizon stochastic problem (32) and apply the first step in the resulting affine control.

The results are shown in Table III. The optimal affine, CE-MPC, and AR-MPC have comparable performances with a very slight edge of MPC methods over the optimal affine controller; AR-MPC has a slight but not statistically significant edge over CE-MPC. (We remind the reader that the online computation required for the CE-MPC and AR-MPC cases is far higher than for the affine controller, in which the optimization is carried out once, offline.) Note that the prescient control law achieves a mean cost of 166.5 with a standard deviation of 54.7. Since the mean cost of the prescient control law is a lower bound on the globally optimal cost, this means that the performance of AR-MPC, CE-MPC, and of the optimal affine controller are all close to optimal.




Fig. 2. Histograms of the total cost incurred using (a) the optimal affine controller, (b) CE-MPC, (c) AR-MPC, and (d) the prescient controller.

Fig. 2 shows the histograms of the total costs achieved by the optimal controller, CE-MPC, AR-MPC, and the prescient controller.

IX. CONCLUSION

We have described a method for effectively computing the globally optimal affine controller for a discrete-time time-varying linear dynamical system operating over a finite horizon, with convex cost function, in a stochastic or worst-case setting. We do this by reducing the problem to a finite-dimensional convex stochastic optimization problem (for the stochastic control problem), or a robust convex optimization problem (in the worst-case setting). While both of these problems can be challenging to solve exactly, there are simple and effective methods, for example based on sampling, that can be used to find good nearly optimal solutions. We have demonstrated our method on three rather different applications, in two cases comparing our method to some other obvious methods.

While the performance achieved by the affine controllers in the given examples (and many others) is very good, we mention again that we are not endorsing affine controllers, or suggesting that they be used instead of other methods, such as CE-MPC. We are merely observing that affine controllers can be effectively computed, and that they appear to give very good control performance. In addition, we make no claim about how suboptimal affine controllers are, or can be. Indeed, this is a very interesting, and we think generally open, research question.

APPENDIX ASSET RETURN MODEL

We model the returns $r(t)$ such that $1 + r(t) = \exp(z(t))$ where $z(t)$ are IID with normal distribution *i.e.*, $z(t) \sim \mathcal{N}(\mu, \Sigma)$. Here $\exp(\cdot)$ is applied entrywise, $\mu \in \mathbf{R}^n$, and $\Sigma \in \mathbf{R}^{n \times n}$.

We consider the case where z is described by a single-factor model, *i.e.*, $z_i = \mu_i + \beta_i g + e_i$, for $i = 1, \dots, n$. Here β_i 's are the factor loadings, g is the (market) factor, and e_i 's represent the residual risks. We take g and e_i 's to be independent Gaussian zero-mean random variables, with respective standard deviations σ_m and σ_i . We can think of σ_m as the market-related standard deviation and σ_i as the firm-specific standard deviation. Under this model, the covariance matrix of z is

$$\Sigma = \sigma_m^2 \beta \beta^T + \text{diag}(\sigma_1^2, \dots, \sigma_n^2)$$

where $\beta = (\beta_1, \dots, \beta_n)$. The mean returns of the assets are

$$\begin{aligned}\mu_i &= \mu_{\text{rf}} + \rho(\mathbf{E}(\beta_i w + v_i)^2)^{1/2} \\ &= \mu_{\text{rf}} + \rho(\beta_i^2 \sigma_m^2 + \sigma_i^2)^{1/2}, \quad i = 1, \dots, n\end{aligned}$$

where μ_{rf} is the risk-free return and ρ is the reward-to-risk ratio. The mean and covariance of $r(t)$ are

$$\begin{aligned}\bar{r} &= \mathbf{E}(r(t)) = \exp(\mu + 0.5 \mathbf{diag}(\Sigma)) - \mathbf{1}, \\ \Sigma_r &= \mathbf{E}((r(t) - \bar{r})(r(t) - \bar{r})^T) \\ &= (\mathbf{1} + \bar{r})(\mathbf{1} + \bar{r})^T \circ (\exp(\Sigma) - \mathbf{1}\mathbf{1}^T)\end{aligned}$$

where \circ indicates the Hadamard product and $\exp(\cdot)$ is applied entrywise.

In Section VII-A, we take the market standard deviation to be $\sigma_m = 10\%$, β_i are chosen uniformly on $[0.3, 1]$ for $i = 2, \dots, n$, and σ_i are chosen uniformly on $[0, 20\%]$ for $i = 2, \dots, n$. We take $\mu_{\text{rf}} = 2\%$, $\rho = 0.4$, and the first asset to be risk-free (*i.e.*, $\beta_1 = 0$, $\sigma_1 = 0$). The four assets have risk (standard deviation) ranging from 0 (for the risk-free asset) to 15.93%, and mean returns ranging from 0 (the risk-free return) to 8.37%. We reorder the assets by increasing risk (and return).

ACKNOWLEDGMENT

The authors thank L. El Ghaoui and U. Topcu for suggesting the problem and for their helpful discussions of it, and J. Primbs for helpful comments.

REFERENCES

- [1] V. Anantharam and C. Desoer, "On the stabilization of nonlinear systems," *IEEE Trans. Autom. Control*, vol. AC-29, no. 6, pp. 569–572, Jun. 1984.
- [2] B. Anderson and J. Moore, *Optimal Control – Linear Quadratic Methods*. Englewood Cliffs, NJ: Prentice-Hall, 1990.
- [3] K. Åström, *Introduction to Stochastic Control Theory*. New York: Dover Publications, 2006.
- [4] R. Bellman, *Dynamic Programming*. New York: Courier Dover Publications, 1957.
- [5] R. Bellman and S. Dreyfus, *Applied Dynamic Programming*. Princeton, NJ: Princeton Univ. Press, 1962.
- [6] A. Bemporad, "Model predictive control design: New trends and tools," in *Proc. 45th IEEE Conf. Decision Control*, 2006, pp. 6678–6683.
- [7] A. Ben-Tal, S. Boyd, and A. Nemirovski, Control of Uncertainty-Affected Discrete Time Linear Systems via Convex Programming. Technical Report Minerva Optimization Center, Technion, Haifa, Israel, 2005 [Online]. Available: http://www.stanford.edu/~boyd/papers/pur_out_control.html
- [8] A. Ben-Tal, S. Boyd, and A. Nemirovski, "Extending scope of robust optimization: Comprehensive robust counterparts of uncertain problems," *Math. Programming*, vol. B, no. 107, pp. 63–89, 2006.
- [9] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski, Robust Optimization 2007 [Online]. Available: www2.isye.gatech.edu/~nemirovs/
- [10] A. Ben-Tal, A. Goryashko, E. Guslitzer, and A. Nemirovski, "Adjustable robust solutions of uncertain linear programs," *Math. Programming*, vol. 99, pp. 351–376, 2004.
- [11] A. Ben-Tal and A. Nemirovski, "Robust convex optimization," *Math. Oper. Res.*, vol. 23, no. 4, pp. 769–805, 1998.
- [12] A. Ben-Tal and A. Nemirovski, "Selected topics in robust convex optimization," *Math. Programming: Series A and B*, vol. 112, no. 1, pp. 125–158, 2007.
- [13] D. Bertsekas, *Neuro-Dynamic Programming*. Nashua, NH: Athena Scientific, 1996.
- [14] D. Bertsekas, *Dynamic Programming and Optimal Control*. Nashua, NH: Athena Scientific, 2005.
- [15] D. Bertsimas, D. Iancu, and P. Parrilo, Optimality of Affine Policies in Multi-Stage Robust Optimization LIDS Tech. Rep. 2809, 2009.
- [16] J. Birge, "Decomposition and partitioning methods for multistage stochastic linear programs," *Oper. Res.*, vol. 33, no. 5, pp. 989–1007, 1985.
- [17] J. Birge and F. Louveaux, "A multicut algorithm for two-stage stochastic linear programs," *Eur. J. Oper. Res.*, vol. 34, pp. 384–392, 1988.
- [18] J. Birge and F. Louveaux, *Introduction to Stochastic Programming*. New York: Springer, 1997.
- [19] R. Bitmead, M. Gevers, and V. Wertz, *Adaptive optimal control: The thinking man's GPC*. Englewood Cliffs, NJ: Prentice-Hall, 1990.
- [20] S. Boyd and C. Barratt, *Linear Controller Design: Limits of Performance*. Englewood Cliffs, NJ: Prentice-Hall, 1991.
- [21] S. Boyd, C. Barratt, and S. Norman, "Linear controller design: Limits of performance via convex optimization," *Proc. IEEE*, vol. 78, no. 3, pp. 529–574, 1990.
- [22] S. Boyd, L. El Ghaoui, E. Feron, and V. Balakrishnan, *Linear Matrix Inequalities in Systems and Control Theory*. Philadelphia, PA: SIAM books, 1994.
- [23] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, 2004.
- [24] G. Calafiore and M. Campi, "On two-stage portfolio allocation problems with affine recourse," in *Proc. 44th IEEE Conf. Decision Control, Eur. Control Conf.*, 2005, pp. 8042–8047.
- [25] G. Calafiore and F. Dabbene, "Probabilistic robust control," in *Proc. Amer. Control Conf.*, 2007, pp. 147–158.
- [26] M. Casey and S. Sen, "The scenario generation algorithm for multi-stage stochastic linear programming," *Math. Oper. Res.*, vol. 30, no. 3, pp. 615–631, 2005.
- [27] X. Chen, M. Sim, P. Sun, and J. Zhang, "A linear decision-based approximation approach to stochastic programming," *Oper. Res.*, vol. 56, no. 2, pp. 344–357, 2008.
- [28] E. Cho, K. Thoney, T. Hodgson, and R. King, "Supply chain planning: Rolling horizon scheduling of multi-factory supply chains," in *Proc. 35th Conf. Winter Simul.: Driving Innovation*, 2003, pp. 1409–1416.
- [29] C. Cutler, "Dynamic Matrix Control: An Optimal Multivariable Control Algorithm with Constraints," Ph.D. dissertation, Univ. Houston, Houston, TX, 1983.
- [30] M. Dahleh and I. Diaz-Bobillo, *Control of Uncertain Systems: A Linear Programming Approach*. Englewood Cliffs, NJ: Prentice-Hall, 1995.
- [31] G. Dantzig and G. Infanger, "Multi-stage stochastic linear programs for portfolio optimization," *Annals Oper. Res.*, vol. 45, pp. 59–76, 1993.
- [32] C. Desoer and C. Lin, "Non-linear unity-feedback systems and Q -parametrization," *Int. J. Control*, vol. 40, no. 1, pp. 37–51, 1984.
- [33] C. Desoer and R. Liu, "Global parametrization of feedback systems," in *Proc. 20th IEEE Conf. Decision Control Symp. Adaptive Processes*, 1981, vol. 20, pp. 859–861.
- [34] N. Elia and M. Dahleh, *Computational Methods for Controller Design*. New York: Springer-Verlag, 1998.
- [35] G. Franklin, J. Powell, and A. Emami-Naeini, *Feedback Control of Dynamic Systems*. Reading, MA: Addison-Wesley, 1994.
- [36] R. Freeman and P. Kokotović, "Optimal nonlinear controllers for feedback linearizable systems," in *Proc. Amer. Control Conf.*, Seattle, WA, 1995, pp. 2722–2726.
- [37] R. Freeman and J. Primbs, "Control Lyapunov functions: New ideas from an old source," in *Proc. 35th IEEE Conf. Decision Control*, Kobe, Japan, 1996, pp. 3926–3931.
- [38] H. Gassmann, "MSLiP: A computer code for the multistage stochastic linear programming problem," *Math. Programming*, vol. 47, pp. 407–423, 1990.
- [39] J. Geromel, C. de Souza, and R. Skelton, "Static output feedback controllers: Stability and convexity," *IEEE Trans. Autom. Control*, vol. 43, no. 1, pp. 120–125, Jan. 1998.
- [40] J. Geromel, P. Peres, and S. Souza, "Convex analysis of output feedback control problems: Robust stability and performance," *IEEE Trans. Autom. Control*, vol. 41, no. 7, pp. 997–1003, Jul. 1996.
- [41] P. Goulart and E. Kerrigan, "Output feedback receding horizon control of constrained systems," *Int. J. Control*, vol. 80, no. 1, pp. 8–20, 2007.
- [42] P. Goulart and E. Kerrigan, "Input-to-state stability of robust receding horizon control with an expected value cost," *Automatica*, vol. 44, no. 4, pp. 1171–1174, 2008.
- [43] P. Goulart, E. Kerrigan, and J. Maciejowski, "Optimization over state feedback policies for robust control with constraints," *Automatica*, vol. 42, no. 4, pp. 523–533, 2006.
- [44] A. Gupta, M. Pál, R. Ravi, and A. Sinha, "What about Wednesday? Approximation algorithms for multistage stochastic optimization," *Approximation, Randomization Combinatorial Optim.*, vol. 3624, pp. 86–98, 2005.

- [45] C. Gustafson and C. Desoer, "Controller design for linear multivariable feedback systems with stable plants, using optimization with inequality constraints," *Int. J. Control*, vol. 37, no. 5, pp. 881–907, 1983.
- [46] A. Hansson, S. Boyd, L. Vandenberghe, and M. Lobo, Optimal Linear Static Control with Moment and Yield Objectives 1997 [Online]. Available: <http://www.stanford.edu/~boyd/papers/yield.html>
- [47] H. Hindi, B. Hassibi, and S. Boyd, "Multi-objective H_2/H_∞ -optimal control via finite dimensional Q -parametrization and linear matrix inequalities," in *Proc. Amer. Control Conf.*, 1998, vol. 5, pp. 3244–3249.
- [48] P. Kall and J. Mayer, *Stochastic Linear Programming: Models, Theory, and Computation*. New York: Springer, 2005.
- [49] V. Kucera, *Discrete Linear Control: The Polynomial Equation Approach*. New York: Wiley, 1979.
- [50] W. Kwon and S. Han, *Receding Horizon Control*. New York: Springer-Verlag, 2005.
- [51] J. Löfberg, "Approximations of closed-loop minimax MPC," in *Proc. 42nd IEEE Conf. Decision Control*, 2003, vol. 2, pp. 1438–1442.
- [52] J. Löfberg, "Minimax Approaches to Robust Model Predictive Control," Ph.D. dissertation, Linköping Univ., Linköping, Sweden, 2003.
- [53] S. Meyn, *Control Techniques for Complex Networks*. Cambridge, U.K.: Cambridge Univ. Press, 2006.
- [54] A. Mutapic and S. Boyd, Cutting-Set Methods for Robust Convex Optimization with Pessimizing Oracles. Manuscript 2008 [Online]. Available: http://www.stanford.edu/~boyd/papers/prac_robust.html
- [55] W. Powell, *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. New York: Wiley, 2007.
- [56] S. Prajna, P. Parrilo, and A. Rantzer, "Nonlinear control synthesis by convex optimization," *IEEE Trans. Autom. Control*, vol. 49, no. 2, pp. 310–314, Feb. 2004.
- [57] A. Prekopa, *Stochastic Programming*. Dordrecht, The Netherlands: Kluwer Academic Publishers, 1995.
- [58] J. Primbs, Dynamic Hedging of Basket Options Under Proportional Transaction Costs Using Receding Horizon Control. Manuscript Tech. Rep., 2007 [Online]. Available: <http://www.stanford.edu/~japrimbs/PapersNew.htm>
- [59] J. Primbs and C. Sung, "Stochastic receding horizon control of constrained linear systems with state and control multiplicative noise," *IEEE Trans. Autom. Control*, submitted for publication.
- [60] M. Rotkowitz and S. Lall, "Affine controller parametrization for decentralized control over Banach spaces," *IEEE Trans. Autom. Control*, vol. 51, no. 9, pp. 1497–1500, Sep. 2006.
- [61] E. Sontag, "A Lyapunov-line characterization of asymptotic controllability," *SIAM J. Control Optim.*, vol. 21, no. 3, pp. 462–471, 1983.
- [62] E. Sontag, "A 'universal' construction of Arstein's theorem on nonlinear stabilization," *Syst. Control Lett.*, vol. 13, no. 2, pp. 117–123, 1989.
- [63] R. Stengel, *Stochastic Optimal Control*. New York: Wiley, 1986.
- [64] C. Sung, "Applications of Modern Control Theory in Portfolio Optimization," Ph.D. dissertation, Stanford Univ., Stanford, CA, 2006.
- [65] C. Swamy and D. Shmoys, "Sampling-based approximation algorithms for multi-stage stochastic optimization," in *Proc. 46th Annu. IEEE Symp. Foundations Comp. Sci.*, 2005, pp. 357–366.
- [66] K. Talluri and G. Van Ryzin, *The Theory and Practice of Revenue Management*. New York: Springer, 2005.
- [67] U. Topcu, G. Calafiore, and L. El Ghaoui, "Multistage investments with recourse: A single-asset case with transaction costs," in *Proc. Conf. Decision Control (CDC'08)*, 2008, pp. 2398–2403.
- [68] D. van Hessem and O. Bosgra, "Closed-loop stochastic dynamic process optimization under input and state constraints," in *Proc. Amer. Control Conf.*, Anchorage, AK, 2002, pp. 2023–2028.
- [69] D. van Hessem and O. Bosgra, "A conic reformulation of model predictive control including bounded and stochastic disturbances and input constraints," in *Proc. Conf. Decision Control*, Las Vegas, NV, 2002, pp. 4643–4648.
- [70] D. van Hessem and O. Bosgra, "A full solution to the constrained stochastic closed-loop MPC problem via state and innovations feedback and its receding horizon implementation," in *Proc. Conf. Decision Control*, Maui, HI, 2003, pp. 929–934.
- [71] D. van Hessem and O. Bosgra, "Closed-loop stochastic model predictive control in a receding horizon implementation on a continuous polymerization reactor example," in *Proc. Amer. Control Conf.*, Boston, MA, 2004, pp. 914–919.
- [72] R. Van Slyke and R. Wets, "L-shaped linear programs with applications to optimal control and stochastic linear programming," *SIAM J. Appl. Math.*, vol. 17, pp. 638–663, 1969.
- [73] M. Vidyasagar, *Control System Synthesis: A Factorization Approach*. Cambridge, MA: MIT Press, 1985.
- [74] C. Wang, C. Ong, and M. Sim, "Model predictive control using affine disturbance feedback for constrained linear system," in *Proc. 46th IEEE Conf. Decision Control*, 2007, pp. 1275–1280.
- [75] P. Whittle, *Optimization Over Time: Dynamic Programming and Stochastic Control*. New York: Wiley, 1982.
- [76] D. Youla, H. Jabr, and J. Bongiorno, "Modern Wiener-Hopf design of optimal controllers-Part II: The multivariable case," *IEEE Trans. Autom. Control*, vol. AC-21, no. 3, pp. 319–338, Jun. 1976.
- [77] D. Youla, J. Bongiorno, and H. Jabr, "Modern Wiener-Hopf design of optimal controllers-Part I: The single-input-output case," *IEEE Trans. Autom. Control*, vol. AC-21, no. 1, pp. 3–13, Feb. 1976.
- [78] K. Zhou, J. Doyle, and K. Glover, *Robust and Optimal Control*. Englewood Cliffs, NJ: Prentice-Hall, 1996.



Joëlle Skaf (M'08) received the B.Eng. degree in computer and communications engineering from the American University of Beirut in 2003, and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 2005 and 2009, respectively.

She is a Software Engineer at Google. Her current interests include convex optimization and its applications in control theory, machine learning, and computational finance.



Stephen P. Boyd (F'99) received the A.B. degree in mathematics (with highest honors) from Harvard University, Cambridge, MA, in 1980 and the Ph.D. degree in electrical engineering and computer science from the University of California, Berkeley, in 1985.

He is the Samsung Professor of Engineering, in the Information Systems Laboratory, Electrical Engineering Department, Stanford University, Stanford, CA. He is the coauthor of *Linear Controller Design: Limits of Performance* (1991), *Linear Matrix Inequalities in System and Control Theory*

(1994), and *Convex Optimization* (2004). His current interests include convex programming applications in control, signal processing, and circuit design.

Dr. Boyd received the ONR Young Investigator Award, the Presidential Young Investigator Award, the 1992 AACC Donald P. Eckman Award, the Perrin Award for Outstanding Undergraduate Teaching in the School of Engineering, the ASSU Graduate Teaching Award, and the AACC Ragazzini Education Award in 2003. He is a Distinguished Lecturer of the IEEE Control Systems Society.