

A Hierarchy of Near-Optimal Policies for Multi-stage Adaptive Optimization

Dimitris Bertsimas ^{*} Dan A. Iancu[†] Pablo A. Parrilo [‡]

June 15, 2009

Abstract

In this paper, we propose a new tractable framework for dealing with multi-stage decision problems affected by uncertainty, applicable to robust optimization and stochastic programming. We introduce a hierarchy of polynomial disturbance-feedback control policies, and show how these can be computed by solving a single semidefinite programming problem. The approach yields a hierarchy parameterized by a single variable (the degree of the polynomial policies), which controls the trade-off between the quality of the objective function value and the computational requirements. We evaluate our framework in the context of two classical inventory management applications, in which very strong numerical performance is exhibited, at relatively modest computational expense.

1 Introduction

Multi-stage optimization problems under uncertainty are prevalent in numerous fields of engineering, economics, finance, and have elicited interest on both a theoretical and a practical level from diverse research communities. Among the most established methodologies for dealing with such problems are dynamic programming (DP) (Bertsekas [2001]), stochastic programming (Birge and Louveaux [2000]), robust control (Zhou and Doyle [1998], Dullerud and Paganini [2005]), and, more recently, robust optimization (see Kerrigan and Maciejowski [2003], Ben-Tal et al. [2005a, 2006], Bertsimas et al. [2009] and references therein).

With a properly defined notion of the *state* of the dynamical system at time k , \mathbf{x}_k , and the controls available to the decision maker, \mathbf{u}_k , one can resort to the Bellman optimality principle of DP (Bertsekas [2001]), to compute optimal policies, $\mathbf{u}_k^*(\mathbf{x}_k)$, and optimal value functions, $J_k^*(\mathbf{x}_k)$. Although DP is a powerful technique as to the theoretical characterization of the optimal policies, it is plagued by the well-known *curse of dimensionality*, in that the complexity of the underlying recursive equations grows quickly with the size of the state-space, rendering the approach ill suited to the computation of actual policy parameters. Therefore, in practice, one would typically solve the recursions numerically (e.g., by multi-parametric programming Bemporad et al. [2000, 2002, 2003]), or resort to approximations, such as approximate DP (de Farias and Van Roy [2003, 2004]), stochastic approximation (Asmussen and Glynn [2007]), simulation based optimization (Marbach

^{*}Sloan School of Management and Operations Research Center, Massachusetts Institute of Technology, 77 Massachusetts Avenue, E40-147, Cambridge, MA 02139, USA. Email: dbertsim@mit.edu.

[†]Operations Research Center, Massachusetts Institute of Technology, 77 Massachusetts Avenue, E40-130, Cambridge, MA 02139, USA. Email: daniancu@mit.edu.

[‡]Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, 77 Massachusetts Avenue, 32D-726, Cambridge, MA 02139, USA. Email: parrilo@mit.edu.

and Tsitsiklis [2001]), and others. Some of the approximations also come with performance guarantees in terms of the objective value in the problem, and many ongoing research efforts are placed on characterizing the sub-optimality gaps resulting from specific classes of policies.

An alternative approach, originally proposed in the stochastic programming community (see Birge and Louveaux [2000], Gatska and Wets [1974] and references therein), is to consider control policies that are parametrized directly in the sequence of observed uncertainties, and typically referred to as *recourse decision rules*. For the case of linear constraints on the controls, with uncertainties regarded as random variables having bounded support and known distributions, and the goal of minimizing an expected piece-wise quadratic, convex cost, the authors in (Gatska and Wets [1974]) show that piece-wise affine decision rules are optimal, but pessimistically conclude that computing the actual parameterization is usually an “impossible task” (for a precise quantification of that statement, see Dyer and Stougie [2006] and Nemirovski and Shapiro [2005]).

Disturbance-feedback parameterizations have recently been used by researchers in robust control and robust optimization (see Löfberg [2003], Kerrigan and Maciejowski [2003, 2004], Goulart and Kerrigan [2005], Ben-Tal et al. [2004, 2005a, 2006], Bertsimas and Brown [2007], Skaf and Boyd [2008a,b], and references therein). In most of the papers, the authors restrict attention to the case of *affine policies*, and show how reformulations can be done that allow the computation of the policy parameters by solving convex optimization problems, which vary from linear and quadratic (e.g. Ben-Tal et al. [2005a], Kerrigan and Maciejowski [2004]), to second-order conic and semidefinite programs (e.g. Löfberg [2003], Ben-Tal et al. [2005a], Bertsimas and Brown [2007], Skaf and Boyd [2008a]). Some of the first steps towards analyzing the properties of disturbance-affine policies were taken in Kerrigan and Maciejowski [2004] and Goulart and Kerrigan [2005], where it was shown that, under suitable conditions, the resulting parametrization has certain desirable system theoretic properties (stability and robust invariance), and that the class of affine disturbance feedback policies is equivalent to the class of affine state feedback policies with memory of prior states, thus subsuming the well-known open-loop and pre-stabilizing control policies.

With the exception of a few classical cases, such as linear quadratic Gaussian or linear exponential quadratic Gaussian¹, characterizing the performance of affine policies in terms of objective function value is typically very hard. The only result in a constrained, robust setting that the authors are aware of is our recent paper Bertsimas et al. [2009], in which it is shown that, in the case of one-dimensional systems, with independent state and control constraints ($L_k \leq u_k \leq U_k, L_k^x \leq x_k \leq U_k^x$), linear control costs and *any* convex state costs, disturbance-affine policies *are*, in fact, optimal, and can be found efficiently. As a downside, the same paper presents simple examples of multi-dimensional systems where affine policies are sub-optimal.

In fact, in most applications, the restriction to the affine case is done for purposes of tractability, and almost invariably results in loss of performance (see the remarks at the end of Nemirovski and Shapiro [2005]), with the optimality gap being sometimes very large. In an attempt to address this problem, recent work has considered parameterizations that are affine in a new set of variables, derived by lifting the original uncertainties into a higher dimensional space. For example, the authors in Chen and Zhang [2009], Chen et al. [2008] suggest using so-called *segregated linear decision rules*, which are affine parameterizations in the positive and negative parts of the original uncertainties. Such policies provide more flexibility, and their computation (for two-stage decision problems in a robust setting) requires the same complexity as that needed for a set of affine policies in the original variables. Another example following similar ideas is Chatterjee et al. [2009], where the authors consider arbitrary functional forms of the disturbances, and show how,

¹These refer to problems that are unconstrained, with Gaussian disturbances, and the goal of minimizing expected costs that are quadratic or exponential of a quadratic, respectively. For these, the optimal policies are affine in the states - see Bertsekas [2001] and references therein.

for specific types of p -norm constraints on the controls, the problems of finding the coefficients of the parameterizations can be relaxed into convex optimization problems. A similar approach is taken in Skaf and Boyd [2008b], where the authors also consider arbitrary functional forms for the policies, and show how, for a problem with convex state-control constraints and convex costs, such policies can be found by convex optimization, combined with Monte-Carlo sampling (to enforce constraint satisfaction). The main drawback of the above approaches is that the *right* choice of functional form for the decision rules is rarely obvious, and there is no systematic way to influence the trade-off between the performance of the resulting policies and the computational complexity required to obtain them, rendering the frameworks ill-suited for general multi-stage dynamical systems, involving complicated constraints on both states and controls.

The goal of our current paper is to introduce a new framework for modeling and (approximately) solving such multi-stage dynamical problems. While we restrict attention mainly to the robust, mini-max objective setting, our ideas can be extended to deal with stochastic problems, in which the uncertainties are random variables with known, bounded support and distribution that is either fully or partially known² (see Section 3.4 for a discussion). Our main contributions are summarized below:

- We introduce a natural extension of the aforementioned affine decision rules, by considering control policies that depend *polynomially* on the observed disturbances. For a fixed polynomial degree d , we develop a convex reformulation of the constraints and objective of the problem, using Sums-Of-Squares (SOS) techniques. In the resulting framework, polynomial policies of degree d can be computed by solving a single semidefinite programming problem (SDP), which, for a fixed precision, can be done in polynomial time Vandenberghe and Boyd [1996]. Our approach is advantageous from a modelling perspective, since it places little burden on the end user (the only choice is the polynomial degree d), while at the same time providing a lever for directly controlling the trade-off between performance and computation (higher d translates into policies with better objectives, obtained at the cost of solving larger SDPs).
- To test our polynomial framework, we consider two classical problems arising in inventory management (single echelon with cumulative order constraints, and serial supply chain with lead-times), and compare the performance of affine, quadratic and cubic control policies. The results obtained are very encouraging - in particular, for all problem instances considered, quadratic policies considerably improve over affine policies (typically by a factor of 2 or 3), while cubic policies essentially close the optimality gap (the relative gap in *all simulations* is less than 1%, with a median gap of less than 0.01%).

The paper is organized as follows. Section 2 presents the mathematical formulation of the problem, briefly discusses relevant solution techniques in the literature, and introduces our framework. Section 3, which is the main body of the paper, first shows how to formulate and solve the problem of searching for the optimal polynomial policy of fixed degree, and then discusses the specific case of polytopic uncertainties. Section 3.4 also elaborates on immediate extensions of the framework to more general multi-stage decision problems. Section 4 translates two classical problems from inventory management into our framework, and Section 5 presents our computational results, exhibiting the strong performance of polynomial policies. Section 6 concludes the paper and suggests directions of future research.

²In the latter case, the cost would correspond to the worst-case distribution consistent with the partial information

1.1 Notation

Throughout the rest of the paper, we denote scalar quantities by lowercase, non-bold face symbols (e.g. $x \in \mathbb{R}, k \in \mathbb{N}$), vector quantities by lowercase, boldface symbols (e.g. $\mathbf{x} \in \mathbb{R}^n, n > 1$), and matrices by uppercase symbols (e.g. $A \in \mathbb{R}^{n \times n}, n > 1$). Also, in order to avoid transposing vectors several times, we use the operator \cdot to denote vertical vector concatenation, e.g. with $\mathbf{x} = (x_1, \dots, x_n) \in \mathbb{R}^n$ and $\mathbf{y} = (y_1, \dots, y_m) \in \mathbb{R}^m$, we write $(\mathbf{x}, \mathbf{y}) \stackrel{\text{def}}{=} (x_1, \dots, x_n, y_1, \dots, y_m) \in \mathbb{R}^{m+n}$.

We refer to quantities specific to time-period k by either including the index in parenthesis, e.g. $\mathbf{x}(k)$, $J^*(k, \mathbf{x}(k))$, or by using an appropriate subscript, e.g. \mathbf{x}_k , $J_k^*(\mathbf{x}_k)$. When referring to the j -th component of a vector at time k , we always use the parenthesis notation for time, and subscript for j , e.g., $\mathbf{x}_j(k)$.

With $\mathbf{x} = (x_1, \dots, x_n)$, we denote by $\mathbb{R}[\mathbf{x}]$ the ring of polynomials in variables x_1, \dots, x_n , and by $\mathcal{P}_d[\mathbf{x}]$ the \mathbb{R} -vector space of polynomials in x_1, \dots, x_n , with degree at most d . We also let

$$\mathcal{B}_d(\mathbf{x}) \stackrel{\text{def}}{=} \left(1, x_1, x_2, \dots, x_n, x_1^2, x_1x_2, \dots, x_1x_n, x_2^2, x_2x_3, \dots, x_n^d\right) \quad (1)$$

be the canonical basis of $\mathcal{P}_d[\mathbf{x}]$, and $s(d) \stackrel{\text{def}}{=} \binom{n+d}{d}$ be its dimension. Any polynomial $f \in \mathcal{P}_d[\mathbf{x}]$ is written as a finite linear combination of monomials,

$$p(\mathbf{x}) = p(x_1, \dots, x_n) = \sum_{\alpha \in \mathbb{N}^n} p_\alpha \mathbf{x}^\alpha = \mathbf{p}^T \mathcal{B}_d(\mathbf{x}), \quad (2)$$

where $\mathbf{x}^\alpha \stackrel{\text{def}}{=} x_1^{\alpha_1} x_2^{\alpha_2} \dots x_n^{\alpha_n}$, and the sum is taken over all n -tuples $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_n) \in \mathbb{N}^n$ satisfying $\sum_{i=1}^n \alpha_i \leq d$. In the expression above, $\mathbf{p} = (p_\alpha) \in \mathbb{R}^{s(d)}$ is the vector of coefficients of $p(\mathbf{x})$ in the basis (1).

2 Problem Description

We consider the following discrete-time, linear dynamical system

$$\mathbf{x}(k+1) = A(k) \mathbf{x}(k) + B(k) \mathbf{u}(k) + \mathbf{w}(k), \quad (3)$$

over a finite planning horizon, $k = 0, \dots, T-1$. The variables $\mathbf{x}(k) \in \mathbb{R}^n$ represent the state, and the controls $\mathbf{u}(k) \in \mathbb{R}^{n_u}$ denote actions taken by the decision maker. $A(k)$ and $B(k)$ are matrices of appropriate dimensions, describing the evolution of the system, and the initial state, $\mathbf{x}(0)$, is assumed known. The system is affected by unknown external disturbances, $\mathbf{w}(k)$, which are assumed to lie in a given compact, basic semialgebraic set,

$$\mathcal{W}_k \stackrel{\text{def}}{=} \{\mathbf{w}(k) \in \mathbb{R}^{n_w} : g_j(\mathbf{w}(k)) \geq 0, j = 1, \dots, m(k)\}, \quad (4)$$

where $g_j \in \mathbb{R}[\mathbf{w}]$ are multivariate polynomials depending on the vector of uncertainties at time k , $\mathbf{w}(k)$. We note that this formulation captures many uncertainty sets of interest, such as polytopic (all g_j affine), p -norms, ellipsoids, and intersections thereof. For simplicity, we omit pre-multiplying $\mathbf{w}(k)$ by a matrix $C(k)$ in (3), since such an evolution could be recast in the current formulation by defining a new uncertainty, $\tilde{\mathbf{w}}(k) = C(k)\mathbf{w}(k)$, evolving in a suitably adjusted set $\tilde{\mathcal{W}}_k$.

We assume that the dynamic evolution is constrained by a set of linear inequalities,

$$\begin{cases} E_x(k) \mathbf{x}(k) + E_u(k) \mathbf{u}(k) \leq \mathbf{f}(k), & k = 0, \dots, T-1, \\ E_x(T) \mathbf{x}(T) \leq \mathbf{f}(T), \end{cases} \quad (5)$$

and the system incurs penalties that are piece-wise affine and convex in the states and controls:

$$h_k(\mathbf{x}_k, \mathbf{u}_k) = \max_{i=1, \dots, r(k)} \left[c_0(k, i) + \mathbf{c}_x(k, i)^T \mathbf{x}(k) + \mathbf{c}_u(k, i)^T \mathbf{u}(k) \right]. \quad (6)$$

The goal is to find non-anticipatory control policies $\mathbf{u}_0, \mathbf{u}_1, \dots, \mathbf{u}_{T-1}$ that minimize the cost incurred by the system in the worst-case scenario. In other words, the problem we seek to solve can be formulated compactly as follows:

$$\min_{\mathbf{u}_0} \left[h_0(\mathbf{x}_0, \mathbf{u}_0) + \max_{\mathbf{w}_0} \min_{\mathbf{u}_1} \left[h_1(\mathbf{x}_1, \mathbf{u}_1) + \dots + \min_{\mathbf{u}_{T-1}} \left[h_{T-1}(\mathbf{x}_{T-1}, \mathbf{u}_{T-1}) + \max_{\mathbf{w}_{T-1}} h_T(\mathbf{x}_T) \right] \dots \right] \right] \quad (7a)$$

$$\text{s.t. } \mathbf{x}_{k+1} = A_k \mathbf{x}_k + B_k \mathbf{u}_k + \mathbf{w}_k, \quad \forall k \in \{0, \dots, T-1\}, \quad (7b)$$

$$(P) \quad E_x(k) \mathbf{x}_k + E_u(k) \mathbf{u}_k \leq \mathbf{f}_k, \quad \forall k \in \{0, \dots, T-1\}, \quad (7c)$$

$$E_x(T) \mathbf{x}_T \leq \mathbf{f}_T. \quad (7d)$$

As already mentioned, the control actions \mathbf{u}_k do not have to be decided entirely at time period $k = 0$, i.e., (P) does not have to be solved as an open-loop problem. Rather, \mathbf{u}_k is allowed to depend on the information set available³ at time k , resulting in control policies $\mathbf{u}_k : \mathcal{F}_k \rightarrow \mathbb{R}^{n_u}$, where \mathcal{F}_k consists of past states, controls and disturbances, $\mathcal{F}_k = \{\mathbf{x}_t\}_{0 \leq t \leq k} \cup \{\mathbf{u}_t\}_{0 \leq t < k} \cup \{\mathbf{w}_t\}_{0 \leq t < k}$.

While \mathcal{F}_k is a large (expanding with k) set, the state \mathbf{x}_k represents sufficient information for taking optimal decisions at time k . Thus, with control policies depending on the states, one can resort to the Bellman optimality principle of Dynamic Programming (DP) Bertsekas [2001], to compute optimal policies, $\mathbf{u}_k^*(\mathbf{x}_k)$, and optimal value functions, $J_k^*(\mathbf{x}_k)$. As suggested in the introduction, the approach is limited due to the *curse of dimensionality*, so that, in practice, one typically resorts to approximate scheme for computing suboptimal, state-dependent policies de Farias and Van Roy [2003, 2004], Asmussen and Glynn [2007], Marbach and Tsitsiklis [2001].

In this paper, we take a slightly different approach, and consider instead policies parametrized directly in the observed uncertainties,

$$\mathbf{u}_k : \mathcal{W}_0 \times \mathcal{W}_1 \times \dots \times \mathcal{W}_{k-1} \rightarrow \mathbb{R}^{n_u}. \quad (8)$$

In this context, with (7b) used to express the dependency of states \mathbf{x}_k on past uncertainties, one typically requires that the state-control constraints (7c), (7d) should be obeyed *robustly* (or *almost surely*, in a stochastic setting), i.e., for any possible realization of the uncertainties. While alternative formulations are possible (e.g., enforcing solutions that obey the constraints with high probability, leading to so-called *chance constraints* Birge and Louveaux [2000]), we do not pursue them further in the current paper, and focus on the former sense of constraint satisfaction.

As already mentioned, this approach has been explored before in the literature, in both the stochastic and robust frameworks (Birge and Louveaux [2000], Gartska and Wets [1974], Löfberg [2003], Kerrigan and Maciejowski [2003, 2004], Goulart and Kerrigan [2005], Ben-Tal et al. [2004, 2005a, 2006], Bertsimas and Brown [2007], Skaf and Boyd [2008a]). The typical restriction to the sub-class of affine policies, done for purposes of tractability, almost invariably results in loss of performance Nemirovski and Shapiro [2005], with the gap being sometimes very large.

To illustrate this effect, we introduce the following simple example⁴, motivated by a similar case in Chen and Zhang [2009]:

³More formally, the decision process u_k is adapted to the filtration generated by past values of the disturbances and controls.

⁴We note that this example can be easily cast as an instance of Problem (P). We opt for the simpler notation to keep the ideas clear.

Example 1. Consider a two-stage problem, where $\mathbf{w} \in \mathcal{W}$ is the uncertainty, with $\mathcal{W} = \{\mathbf{w} \in \mathbb{R}^N : \|\mathbf{w}\|_2 \leq 1\}$, $x \in \mathbb{R}$ is a first-stage decision (taken before \mathbf{w} is revealed), and $\mathbf{u} \in \mathbb{R}^N$ is a second-stage decision (allowed to depend on \mathbf{w}). We would like to solve the following optimization:

$$\begin{aligned} & \underset{x, \mathbf{u}(\mathbf{w})}{\text{minimize}} && x \\ & \text{such that} && x \geq \sum_{i=1}^N u_i, \quad \forall \mathbf{w} \in \mathcal{W}, \\ & && u_i \geq w_i^2, \quad \forall \mathbf{w} \in \mathcal{W}. \end{aligned} \tag{9}$$

It can be easily shown (see Lemma 1 in Section 7.1) that the optimal objective in Problem (9) is 1, corresponding to $u_i(\mathbf{w}) = w_i^2$, while the best objective achievable under *affine* policies $\mathbf{u}(\mathbf{w})$ is N , for $u_i(\mathbf{w}) = 1, \forall i$. In particular, this simple example shows that the optimality gap resulting from the use of affine policies can be made arbitrarily large (as the problem size increases).

Motivated by these facts, in the current paper we explore the performance of a more general class of disturbance-feedback control laws, namely policies that are *polynomial* in past-observed uncertainties. More precisely, for a specified degree d , and with $\boldsymbol{\xi}_k$ denoting the vector of all disturbances in \mathcal{F}_k ,

$$\boldsymbol{\xi}_k \stackrel{\text{def}}{=} [\mathbf{w}_0, \mathbf{w}_1, \dots, \mathbf{w}_{k-1}] \in \mathbb{R}^{k \cdot n_w}, \tag{10}$$

we consider a control law at time k in which every component is a polynomial of degree at most d in variables $\boldsymbol{\xi}_k$, i.e., $u_j(k, \boldsymbol{\xi}_k) \in \mathcal{P}_d[\boldsymbol{\xi}_k]$, and thus:

$$\mathbf{u}_k(\boldsymbol{\xi}_k) = L_k \mathcal{B}_d(\boldsymbol{\xi}_k), \tag{11}$$

where $\mathcal{B}_d(\boldsymbol{\xi}_k)$ is the canonical basis of $\mathcal{P}_d[\boldsymbol{\xi}_k]$, given by (1). The new decision variables become the matrices of coefficients $L_k \in \mathbb{R}^{n_u \cdot s(d)}$, $k = 0, \dots, T-1$, where $s(d) = \binom{k \cdot n_w + d}{d}$ is the dimension of $\mathcal{P}_d[\boldsymbol{\xi}_k]$. Therefore, with a fixed degree d , the number of decision variables remains polynomially bounded in the size of the problem input, T, n_u, n_w .

This class of policies constitutes a natural extension of the disturbance-affine control laws, i.e., the case $d = 1$. Furthermore, with sufficiently large degree, one can expect the performance of the polynomial policies to become near-optimal (recall that, by the Stone-Weierstrass Theorem Rudin [1976], any continuous function on a compact set can be approximated as closely as desired by polynomial functions). The main drawback of the approach is that searching over arbitrary polynomial policies typically results in non-convex optimization problems. To address this issue, in the next section, we develop a tractable, convex reformulation of the problem based on Sums-Of-Squares (SOS) techniques Parrilo [2000, 2003], Lasserre [2001].

3 Polynomial Policies and Convex Reformulations Using Sums-Of-Squares

3.1 Reformulating the Constraints

With polynomial control policies of the form (11), note that a typical state-control constraint (7c, 7d) in program (P) can now be written as:

$$p(\boldsymbol{\xi}) \geq 0, \quad \forall \boldsymbol{\xi} \in \mathcal{W}_0 \times \dots \times \mathcal{W}_{k-1}, \tag{12}$$

where $\boldsymbol{\xi} \in \mathbb{R}^{k \cdot n_w}$ is given by (10), and $p(\boldsymbol{\xi})$ is a polynomial in variables $\xi_1, \xi_2, \dots, \xi_{k \cdot n_w}$ with degree at most d ,

$$p(\boldsymbol{\xi}) = \mathbf{p}^T \mathcal{B}_d(\boldsymbol{\xi}),$$

whose coefficients p_i are affine combinations of the decision variables L_t , $0 \leq t \leq k$. It is easy to see that constraint (12) can be rewritten equivalently as

$$p(\boldsymbol{\xi}) \geq 0, \quad \forall \boldsymbol{\xi} \in \Xi \stackrel{\text{def}}{=} \{\boldsymbol{\xi} \in \mathbb{R}^{k \cdot n_w} : g_j(\boldsymbol{\xi}) \geq 0, j = 1, \dots, m\}, \quad (13)$$

where $\{g_j\}_{1 \leq j \leq m}$ are all the polynomial functions describing the compact basic semi-algebraic set $\Xi \equiv \mathcal{W}_0 \times \dots \times \mathcal{W}_{k-1}$, immediately derived from (4). In this form, (13) falls in the general class of constraints that require testing polynomial non-negativity on a basic closed, semi-algebraic set, i.e., a set given by a finite number of polynomial equalities and inequalities. To this end, note that a *sufficient* condition for (13) to hold is:

$$p = \sigma_0 + \sum_{j=1}^m \sigma_j g_j, \quad (14)$$

where $\sigma_j \in \mathbb{R}[\boldsymbol{\xi}]$, $j = 0, \dots, m$, are polynomials in the variables $\boldsymbol{\xi}$ which are furthermore *sums of squares* (SOS). This condition translates testing the non-negativity of p on the set Ξ into a system of linear equality constraints on the coefficients of p and σ_j , $j = 0, \dots, m$, and a test whether σ_j are SOS. The main reason why this is valuable is because testing whether a polynomial of fixed degree is SOS is equivalent to solving a semidefinite programming problem (SDP) (refer to Parrilo [2000, 2003], Lasserre [2001] for details), which, for a fixed precision, can be done in polynomial time, by interior point methods Vandenberghe and Boyd [1996].

On first sight, condition (14) might seem overly restrictive. However, it is motivated by recent powerful results in real algebraic geometry (Putinar [2003], Jacobi and Prestel [2001]), which, under mild conditions⁵ on the functions g_j , state that *any* polynomial that is strictly positive on a compact semi-algebraic set Ξ *must* admit a representation of the form (14), where the degrees of the σ_j polynomials are not a priori bounded. In our framework, in order to obtain a tractable formulation, we furthermore restrict these degrees so that the total degree of every product $\sigma_j g_j$ is at most d , the degree of the control policies (11) under consideration. While this requirement is more restrictive, and could, in principle, result in conservative parameter choices, it avoids ad-hoc modeling decisions and has the advantage of keeping a single parameter that is adjustable to the user (the degree d), which directly controls the trade-off between the size of the resulting SDP formulation and the quality of the overall solution. Furthermore, in our numerical simulations, we find that this choice performs very well in practice, and never results in infeasible conditions.

3.2 Reformulating the Objective

To complete the description, we now show how to model the objective (7a). Recall that, in our problem, the costs at every time-step k are convex, piecewise affine functions of the state \mathbf{x}_k and the control \mathbf{u}_k . Therefore, when polynomial control policies of the form (11) are used, a typical stage cost (6) becomes a piecewise polynomial function of past uncertainties, i.e., a maximum of several polynomials. A natural way to bring such a cost into the framework presented before is to

⁵These are readily satisfied when g_j are affine, or can be satisfied by simply appending a redundant constraint that bounds the 2-norm of the vector $\boldsymbol{\xi}$

introduce, for every stage $k = 0, \dots, T$, a polynomial function of past uncertainties, and require it to be an upper-bound on the true (piecewise polynomial) cost.

More precisely, and to fix ideas, consider the stage cost at time k , $h_k(\mathbf{x}_k, \mathbf{u}_k)$. With \mathbf{u}_t , $t = 0, \dots, k$, given by (11), the expression (6) for this cost becomes

$$h_k(\mathbf{x}_k, \mathbf{u}_k) = \max_{i=1, \dots, m_k} p_i(\boldsymbol{\xi}_k),$$

where $\boldsymbol{\xi}_k$ is the vector of observed uncertainties, given by (10), and $p_i \in \mathcal{P}_d[\boldsymbol{\xi}_k]$ is the polynomial

$$p_i(\boldsymbol{\xi}_k) = c_0(k, i) + \mathbf{c}_x(k, i)^T \mathbf{x}_k(\boldsymbol{\xi}_k) + \mathbf{c}_u(k, i)^T \mathbf{u}_k(\boldsymbol{\xi}_k).$$

In this context, we introduce a modified stage cost $\tilde{h}_k \in \mathcal{P}_d[\boldsymbol{\xi}_k]$, which we constrain to satisfy

$$\tilde{h}_k(\boldsymbol{\xi}_k) \geq p_i(\boldsymbol{\xi}_k), \quad \forall \boldsymbol{\xi}_k \in \mathcal{W}_0 \times \dots \times \mathcal{W}_{k-1}, \quad \forall i = 1, \dots, m_k, \quad (15)$$

and we replace the overall cost for Problem (P) with the sum of the modified stage costs. In other words, instead of minimizing the objective (7a), we seek to solve:

$$\begin{aligned} \min \quad & J \\ \text{s.t.} \quad & J \geq \sum_{k=0}^T \tilde{h}_k(\boldsymbol{\xi}_k), \quad \forall \boldsymbol{\xi}_T \in \mathcal{W}_0 \times \dots \times \mathcal{W}_{T-1}, \end{aligned} \quad (16)$$

The advantage of this approach is that, now, constraints (15) and (16) are of the exact same nature as (12), and thus fit into the SOS framework developed earlier. As a result, we can use the same semidefinite programming approach to enforce them, while preserving the tractability of the formulation⁶ and the trade-off between performance and computation delivered by the degree d . The main drawback is that the cost J may, in general, strictly over-bound the optimal cost of Problem (P), due to several reasons:

1. We are replacing the (true) piece-wise polynomial cost h_k with an *upper bound* given by the polynomial cost \tilde{h}_k . Therefore, the optimal value J of problem (16) may, in general, be larger than the true cost corresponding to the respective polynomial policies.
2. All the constraints in the model, namely (15), (16) and the state-control constraints (7c), (7d), are enforced using SOS polynomials with fixed degree (see the discussion in Section 3.1), and this is sufficient, but not necessary.

However, despite these multiple layers of approximation, our numerical experiments, presented in Section 5, suggest that most of the above considerations are second-order effects when compared with the fact that polynomial policies of the form (11), are themselves, in general, suboptimal. In fact, our results suggest that with a modest polynomial degree (3, and sometimes even 2), one can close most of the optimality gap between the SDP formulation and the optimal value of Problem (P).

To summarize, our framework can be presented as the sequence of steps below:

⁶We would need to add $\mathcal{O}(T \cdot (\max_k m_k) \cdot (m+1))$ semidefinite constraints, where m is the number of polynomial inequalities specifying the set $\mathcal{W}_0 \times \dots \times \mathcal{W}_{T-1}$, and each semidefinite matrix has size at most $k \cdot n_w$.

Algorithm 1 Framework for computing optimal polynomial policies of degree d

- 1: Consider polynomial control policies in the disturbances, $\mathbf{u}_k(\boldsymbol{\xi}_k) = L_k \mathcal{B}_d(\boldsymbol{\xi}_k)$.
- 2: Replace a typical stage cost $h_k(\mathbf{x}_k, \mathbf{u}_k) = \max_{i=1, \dots, m_k} p_i(\boldsymbol{\xi}_k)$ with a modified stage cost $\tilde{h}_k \in \mathcal{P}_d[\boldsymbol{\xi}_k]$, constrained to satisfy $\tilde{h}_k(\boldsymbol{\xi}_k) \geq p_i(\boldsymbol{\xi}_k), \forall \boldsymbol{\xi}_k \in \mathcal{W}_0 \times \dots \times \mathcal{W}_{k-1}, \forall i = 1, \dots, m_k$.
- 3: Replace the overall cost for Problem (P) with the sum of the modified stage costs.
- 4: Replace a typical constraint $p(\boldsymbol{\xi}_k) \geq 0, \forall \boldsymbol{\xi}_k \in \{\boldsymbol{\xi} : g_j(\boldsymbol{\xi}) \geq 0, j = 1, \dots, m\}$ (for either state-control or costs) with the requirements:

$$p = \sigma_0 + \sum_{j=1}^m \sigma_j g_j \quad (\text{linear equality constraints on coefficients})$$

$$\sigma_j \text{ SOS, } \deg(\sigma_0) \leq d, \deg(\sigma_j g_j) \leq d, j = 1, \dots, m. \quad (m + 1 \text{ SDP constraints})$$

- 5: Solve the resulting SDP to obtain the coefficients L_k of the policies.
-

3.3 An Exact Value for Polytopic Uncertainties.

In this section, we briefly discuss a specific case of Problem (P), in which the *exact* optimal value of the optimization can be computed by solving a (large) mathematical program. This is particularly useful for benchmarking purposes, since it allows a precise assessment of the polynomial framework's performance (note that the approach presented in the previous section is applicable to the general problem, described in the introduction).

Consider the particular case of polytopic uncertainty sets, i.e., when all the polynomial functions g_j in (4) are actually affine. It can be shown (see Theorem 2 in Bemporad et al. [2003]) that piecewise affine state-feedback policies⁷ $\mathbf{u}_k(\mathbf{x}_k)$ are optimal for the resulting Problem (P), and that the sequence of uncertainties that achieves the min-max value is an extreme point of the uncertainty set, that is, $(\mathbf{w}_0, \mathbf{w}_1, \dots, \mathbf{w}_{T-1}) \in \text{ext}(\mathcal{W}_0) \times \dots \times \text{ext}(\mathcal{W}_{T-1})$. As an immediate corollary of this result, the optimal value for Problem (P), as well as the optimal decision at time $k = 0$ for a fixed initial state $\mathbf{x}_0, \mathbf{u}_0^*(\mathbf{x}_0)$, can be computed by solving the following optimization problem (see Ben-Tal et al. [2005a], Bemporad et al. [2002, 2003] for a proof):

$$\min_{\mathbf{u}_k(\boldsymbol{\xi}_k), z_k(\boldsymbol{\xi}_k), J} J \quad (17a)$$

$$J \geq \sum_{k=0}^T z_k(\boldsymbol{\xi}_k), \quad (17b)$$

$$z_k(\boldsymbol{\xi}_k) \geq h_k(\mathbf{x}_k(\boldsymbol{\xi}_k), \mathbf{u}_k(\boldsymbol{\xi}_k)), \quad k = 0, \dots, T-1, \quad (17c)$$

$$(P)_{\text{ext}} \quad z_T(\boldsymbol{\xi}_T) \geq h_T(\mathbf{x}_T(\boldsymbol{\xi}_T)), \quad (17d)$$

$$\mathbf{x}_{k+1}(\boldsymbol{\xi}_{k+1}) = A_k \mathbf{x}_k(\boldsymbol{\xi}_k) + B_k \mathbf{u}_k(\boldsymbol{\xi}_k) + \mathbf{w}(k), \quad k = 0, \dots, T-1, \quad (17e)$$

$$\mathbf{f}_k \geq E_x(k) \mathbf{x}_k(\boldsymbol{\xi}_k) + E_u(k) \mathbf{u}_k(\boldsymbol{\xi}_k), \quad k = 0, \dots, T-1, \quad (17f)$$

$$\mathbf{f}_T \geq E_x(T) \mathbf{x}_T(\boldsymbol{\xi}_T). \quad (17g)$$

In this formulation, *non-anticipatory control values* $\mathbf{u}_k(\boldsymbol{\xi}_k)$ and corresponding states $\mathbf{x}_k(\boldsymbol{\xi}_k)$ are computed for every vertex of the disturbance set, i.e., for every $\boldsymbol{\xi}_k \in \text{ext}(\mathcal{W}_0) \times \dots \times \text{ext}(\mathcal{W}_{k-1}), k = 0, \dots, T-1$. The variables $z_k(\boldsymbol{\xi}_k)$ are used to model the stage cost at time k , in scenario $\boldsymbol{\xi}_k$. Note

⁷One could also immediately extend the result of Gatska and Wets [1974] to argue that disturbance-feedback policies $\mathbf{u}_k(\mathbf{w}_1, \dots, \mathbf{w}_{k-1})$ are also optimal.

that constraints (17c), (17d) can be immediately rewritten in linear form, since the functions $h_k(\mathbf{x}, \mathbf{u})$, $h_T(\mathbf{x})$ are piece-wise affine and convex in their arguments.

We emphasize that the formulation does not seek to compute an actual *policy* $\mathbf{u}_k^*(\mathbf{x}_k)$, but rather the values that this policy would take (and the associated states and costs) when the uncertainty realizations are restricted to extreme points of the uncertainty set. As such, the variables $\mathbf{u}_k(\boldsymbol{\xi}_k)$, $\mathbf{x}_k(\boldsymbol{\xi}_k)$ and $z_k(\boldsymbol{\xi}_k)$ must also be forced to satisfy a *non-anticipativity* constraint, which is implicitly taken into account when only allowing them to depend on the portion of the extreme sequence available at time k , i.e., $\boldsymbol{\xi}_k$. Due to this coupling constraint, Problem $(P)_{\text{ext}}$ results in a Linear Program which is doubly-exponential in the horizon T , with the number of variables and the number of constraints both proportional to the number of extreme sequences in the uncertainty set, $\mathcal{O}\left(\prod_{k=0}^{T-1} |\text{ext}(\mathcal{W}_k)|\right)$. Therefore, solving $(P)_{\text{ext}}$ is relevant only for small horizons, but is very useful for benchmarking purposes, since it provides the optimal value of the original problem.

We conclude this section by examining a particular example when the uncertainty sets take a simpler form, and polynomial policies (11) are provably optimal. More precisely, we consider the case of scalar uncertainties ($n_w = 1$), and

$$\mathbf{w}(k) \in \mathcal{W}(k) \stackrel{\text{def}}{=} [\underline{w}_k, \overline{w}_k] \subset \mathbb{R}, \quad \forall k = 0, \dots, T-1, \quad (18)$$

known in the literature as *box uncertainty* Ben-Tal and Nemirovski [2002], Ben-Tal et al. [2004]. Under this model, any partial uncertain sequence $\boldsymbol{\xi}_k \stackrel{\text{def}}{=} (w_0, \dots, w_{k-1})$ will be a k -dimensional vector, lying inside the hypercube $\Xi_k \stackrel{\text{def}}{=} \mathcal{W}_0 \times \dots \times \mathcal{W}_{k-1} \subset \mathbb{R}^k$.

Introducing the subclass of *multi-affine* policies⁸ of degree d , given by

$$u_j(k, \boldsymbol{\xi}_k) = \sum_{\boldsymbol{\alpha} \in \{0,1\}^k} \ell_{\boldsymbol{\alpha}} \boldsymbol{\xi}_k^{\boldsymbol{\alpha}}, \quad \text{where} \quad \sum_{i=1}^k \alpha_i \leq d, \quad (19)$$

one can show (see Theorem 1 in the Appendix) that multi-affine policies of degree $T-1$ are, in fact, optimal for Problem (P) . While this theoretical result is of minor practical importance (due to the large degree needed for the policies, which translates into prohibitive computation), it provides motivation for restricting attention to polynomials of smaller degree, as a midway solution that preserves tractability, while delivering high quality objective values.

3.4 Extensions

For completeness, we conclude our discussion by briefly mentioning several modelling extensions that can be readily captured in our framework:

1. Polynomial state-control constraints, i.e., when (5) are replaced with $p_j(\mathbf{x}_k, \mathbf{u}_k) \leq 0$, where $p_j \in \mathcal{P}_q[\mathbf{x}_k, \mathbf{u}_k]$, for some degree q . This could capture very general constraint structures (p -norm, ellipsoidal, and even non-convex!), which might be of interest in specific applications. Of course, the pitfall of the approach is that the size of the resulting SDPs would grow, as the degree of a polynomial in a typical constraint could be as large as $q \cdot d$.
2. Polynomial system dynamics, i.e., $\mathbf{x}_{k+1} = p_k(\mathbf{x}_k, \mathbf{u}_k) + \mathbf{w}_k$, with $p_k \in \mathcal{P}_q[\mathbf{x}_k, \mathbf{u}_k]$. Same observations as above apply.

⁸Note that these are simply polynomial policies of the form (11), involving only square-free monomials, i.e., every monomial, $\boldsymbol{\xi}_k^{\boldsymbol{\alpha}} \stackrel{\text{def}}{=} \prod_{i=1}^k \xi_i^{\alpha_i}$, satisfies the condition $\alpha_i \in \{0, 1\}$.

3. Note that, instead of considering uncertainties as lying in given sets, and adopting a min-max (worst-case) objective, we could accommodate the following modelling assumptions:

- (a) The uncertainties are random variables, with bounded support given by the set $\mathcal{W}_0 \times \mathcal{W}_1 \times \dots \times \mathcal{W}_{T-1}$, and known probability distribution function \mathbb{F} . The goal is to find $\mathbf{u}_0, \dots, \mathbf{u}_{T-1}$ so as to obey the state-control constraints (5) almost surely, and to minimize the expected costs,

$$\min_{\mathbf{u}_0} \left[h_0(\mathbf{x}_0, \mathbf{u}_0) + \mathbb{E}_{\mathbf{w}_0 \sim \mathbb{F}} \min_{\mathbf{u}_1} \left[h_1(\mathbf{x}_1, \mathbf{u}_1) + \dots + \min_{\mathbf{u}_{T-1}} \left[h_{T-1}(\mathbf{x}_{T-1}, \mathbf{u}_{T-1}) + \mathbb{E}_{\mathbf{w}_{T-1} \sim \mathbb{F}} h_T(\mathbf{x}_T) \right] \dots \right] \right]. \quad (20)$$

In this case, since our framework already enforces almost sure (robust) constraint satisfaction, the only potential modifications would be in the reformulation of the objective. Since the distribution of the uncertainties is assumed known, and the support is bounded, the moments exist and can be computed up to any fixed degree d . Therefore, we could preserve the reformulation of state-control constraints and stage-costs in our framework (i.e., Steps 2 and 4), but then proceed to minimize the *expected* sum of the polynomial costs \tilde{h}_k (note that the expected value of a polynomial function of uncertainties can be immediately obtained as a linear function of the moments).

- (b) The uncertainties are random variables, with the same bounded support as above, but unknown distribution function \mathbb{F} , belonging to a given set of distributions, \mathcal{F} . The goal is to find control policies obeying the constraints almost surely, and minimizing the expected costs corresponding to the *worst-case distribution* \mathbb{F} ,

$$\min_{\mathbf{u}_0} \left[h_0(\mathbf{x}_0, \mathbf{u}_0) + \sup_{\mathbb{F} \in \mathcal{F}} \mathbb{E}_{\mathbf{w}_0} \min_{\mathbf{u}_1} \left[h_1(\mathbf{x}_1, \mathbf{u}_1) + \dots + \min_{\mathbf{u}_{T-1}} \left[h_{T-1}(\mathbf{x}_{T-1}, \mathbf{u}_{T-1}) + \sup_{\mathbb{F} \in \mathcal{F}} \mathbb{E}_{\mathbf{w}_{T-1}} h_T(\mathbf{x}_T) \right] \dots \right] \right]. \quad (21)$$

In this case, if partial information (such as the moments of the distribution up to degree d) is available, then the framework in (a) could be applied. Otherwise, if the only information available about \mathbb{F} were the support, then our framework could be applied without modification, but the solution obtained would exactly correspond to the min-max approach, and hence be quite conservative.

While these extensions are certainly worthy of attention, we do not pursue them here, and restrict our tests in the remainder of the paper to the worst-case cost formulation.

4 Examples from Inventory Management

To test the performance of our proposed policies, we consider two problems arising in inventory management.

4.1 Single Echelon with Cumulative Order Constraints

This first example was originally discussed in a robust framework by Ben-Tal et al. [2005b], in the context of a more general model for the problem of negotiating flexible contracts between a retailer

and a supplier in the presence of uncertain orders from customers. We describe a simplified version of the problem, which is sufficient to illustrate the benefit of our approach, and refer the interested reader to Ben-Tal et al. [2005b] for more details.

The setting is the following: consider a single-product, single-echelon, multi-period supply chain, in which inventories are managed periodically over a planning horizon of T periods. The unknown demands w_k from customers arrive at the (unique) echelon, henceforth referred to as the *retailer*, and are satisfied from the on-hand inventory, denoted by x_k at the beginning of period k . The retailer can replenish the inventory by placing orders u_k , at the beginning of each period k , for a cost of c_k per unit of product. These orders are immediately available, i.e., there is no lead-time in the system, but there are capacities on the order size in every period, $L_k \leq u_k \leq U_k$, as well as on the cumulative orders placed in consecutive periods, $\hat{L}_k \leq \sum_{t=0}^k u_t \leq \hat{U}_k$. After the demand w_k is realized, the retailer incurs holding costs $H_{k+1} \cdot \max\{0, x_k + u_k - w_k\}$ for all the amounts of supply stored on her premises, as well as penalties $B_{k+1} \cdot \max\{w_k - x_k - u_k, 0\}$, for any demand that is backlogged.

In the spirit of robust optimization, we assume that the only information available about the demand at time k is that it resides within an interval centered around a *nominal* (mean) demand \bar{d}_k , which results in the uncertainty set $\mathcal{W}_k = \{w_k \in \mathbb{R} : |w_k - \bar{d}_k| \leq \rho \cdot \bar{d}_k\}$, where $\rho \in [0, 1]$ can be interpreted as an *uncertainty level*.

With the objective function to be minimized as the cost resulting in the worst-case scenario, we immediately obtain an instance of our original Problem (P), i.e., a linear system with $n = 2$ states and $n_u = 1$ control, where $x_1(k)$ represents the on-hand inventory at the beginning of time k , and $x_2(k)$ denotes the total amount of orders placed in prior times, $x_2(k) = \sum_{t=0}^{k-1} u(t)$. The dynamics are specified by

$$\begin{aligned} x_1(k+1) &= x_1(k) + u(k) - w(k), \\ x_2(k+1) &= x_2(k) + u(k), \end{aligned}$$

with the constraints

$$\begin{aligned} L_k &\leq u(k) \leq U_k, \\ \hat{L}_k &\leq x_2(k) + u(k) \leq \hat{U}_k, \end{aligned}$$

and the costs

$$\begin{aligned} h_k(\mathbf{x}_k, u_k) &= \max\{c_k u_k + [H_k, 0]^T \mathbf{x}_k, c_k u_k + [-B_k, 0]^T \mathbf{x}_k\}, \\ h_T(\mathbf{x}_T) &= \max\{[H_T, 0]^T \mathbf{x}_T, [-B_T, 0]^T \mathbf{x}_T\}. \end{aligned}$$

We remark that the cumulative order constraints, $\hat{L}_k \leq \sum_{t=0}^k u_t \leq \hat{U}_k$, are needed here, since otherwise, the resulting (one-dimensional) system would fit the theoretical results from Bertsimas et al. [2009], which would imply that polynomial policies of the form (11) and polynomial stage costs of the form (15) are already optimal for degree $d = 1$ (affine). Therefore, testing for higher order polynomial policies would not add any benefit.

4.2 Serial Supply Chain

As a second problem, we consider a serial supply chain, in which there are J echelons, numbered $1, \dots, J$, managed over a planning horizon of T periods by a centralized decision maker. The j -th echelon can hold inventory on its premises, for a per-unit cost of $H_j(k)$ in time period k . In every period, echelon 1 faces the unknown, external demands $w(k)$, which it must satisfy from

the on-hand inventory. Unmet demands can be backlogged, incurring a particular per-unit cost, $B_1(k)$. The j -th echelon can replenish its on-hand inventory by placing orders with the immediate echelon in the upstream, $j + 1$, for a per-unit cost of $c_j(k)$. For simplicity, we assume the orders are received with zero lead-time, and are only constrained to be non-negative, and we assume that the last echelon, J , can replenish inventory from a supplier with infinite capacity.

Following a standard requirement in inventory theory Zipkin [2000], we maintain that, under centralized control, orders placed by echelon j at the beginning of period k cannot be backlogged at echelon $j + 1$, and thus must always be sufficiently small to be satisfiable from on-hand inventory at the beginning⁹ of period k at echelon $j + 1$. As such, instead of referring to orders placed by echelon j to the upstream echelon $j + 1$, we will refer to physical shipments from $j + 1$ to j , in every period.

This problem can be immediately translated into the linear systems framework mentioned before, by introducing the following states, controls, and uncertainties:

- Let $x_j(k)$ denote the local inventory at stage j , at the beginning of period k .
- Let $u_j(k)$ denote the shipment sent in period k from echelon $j + 1$ to echelon j .
- Let the unknown external demands arriving at echelon 1 represent the uncertainties, $w(k)$.

The dynamics of the linear system can then be formulated as

$$\begin{aligned} x_1(k+1) &= x_1(k) + u_1(k) - w(k), & k = 0, \dots, T-1, & \quad (\text{local inventory at echelon 1}) \\ x_j(k+1) &= x_j(k) + u_j(k) - u_{j-1}(k), & j = 2, \dots, J, k = 0, \dots, T-1, & \quad (\text{local inventory at echelon } j > 1) \end{aligned}$$

with the following constraints on the states and controls

$$\begin{aligned} u_j(k) &\geq 0, & j = 1, \dots, J, k = 0, \dots, T-1, & \quad (\text{non-negative shipments}) \\ x_j(k) &\geq u_{j-1}(k), & j = 2, \dots, J, k = 0, \dots, T-1, & \quad (\text{downstream shipment } \leq \text{upstream inventory}) \end{aligned}$$

and the costs

$$\begin{aligned} h_1(k, x_1(k), u_1(k)) &= c_1(k)u_1(k) + \max\{H_1(k)x_1(k), -B_1(k)x_1(k)\}, & k = 0, \dots, T-1 \\ h_1(T, x_1(T)) &= \max\{H_1(T)x_1(T), -B_1(T)x_1(T)\}, \\ h_j(k, x_j(k), u_j(k)) &= c_j(k)u_j(k) + H_j(k)x_j(k), & k = 0, \dots, T-1 \\ h_j(T, x_j(T)) &= H_j(T)x_j(T). \end{aligned}$$

With the same model of uncertainty as before, $\mathcal{W}_k = [\bar{d}_k(1 - \rho), \bar{d}_k(1 + \rho)]$, for some known mean demand \bar{d}_k and uncertainty level $\rho \in [0, 1]$, and the goal to decide shipment quantities $u_j(k)$ so as to minimize the cost in the worst-case scenario, we obtain a different example of Problem (P).

5 Numerical Experiments

In this section, we present numerical simulations testing the performance of polynomial policies in each of the two problems mentioned in Section 4. In order to examine the dependency of our results on the size of the problem, we proceed in the following fashion.

⁹This implies that the order placed by echelon j in period k (to the upstream echelon, $j + 1$) cannot be used to satisfy the order in period k from the downstream echelon, $j - 1$. Technically, this corresponds to an effective lead time of 1 period, and a more appropriate model would redefine the state vector accordingly. We have opted to keep our current formulation for simplicity.

5.1 First Example

For the first model (single echelon with cumulative order constraints), we vary the horizon of the problem from $T = 4$ to $T = 10$, and for every value of T , we:

1. Create 100 problem instances, by randomly generating the cost parameters and the constraints, in which the performance of polynomial policies of degree 1 (affine) is suboptimal.
2. For every such instance, we compute:
 - The optimal cost OPT , by solving the exponential Linear Program $(P)_{\text{ext}}$.
 - The optimal cost \bar{P}_d obtained with polynomial policies of degree $d = 1, 2$, and 3 , respectively, by solving the corresponding associated SDP formulations, as introduced in Section 3.

We also record the relative optimality gap corresponding to each polynomial policy, defined as $(\bar{P}_d - OPT)/OPT$, and the solver time.

3. We compute statistics over the 100 different instances (recording the mean, standard deviation, min, max and median) for the optimality gaps and solver times corresponding to all three polynomial parameterizations.

Table 1 and Table 2 record these statistics for relative gaps and solver times, respectively. The following conclusions can be drawn from the results:

- Policies of higher degree decrease the performance gap considerably. In particular, while affine policies yield an average gap between 2.8% and 3.7% (with a median gap between 2% and 2.7%), quadratic policies reduce both average and median gaps by a factor of 3, and cubic policies essentially close the optimality gap (all gaps are smaller than 1%, with a median gap smaller than 0.01%). To better see this, Figure 1 illustrates the box-plots corresponding to the three policies for a typical case (here, $T = 6$).
- The reductions in the relative gaps are not very sensitive to the horizon, T . Figure 2(a) illustrates this effect for the case of quadratic policies, and similar plots can be drawn for the affine and cubic cases.
- The computational time grows polynomially with the horizon size. While computations for cubic policies are rather expensive, the quadratic case, shown in Figure 2(b), shows promise for scalability - for horizon $T = 10$, the median and average solver times are below 15 seconds.

5.2 Second Example

For the second model (serial supply chain), we fix the problem horizon to $T = 7$, and vary the number of echelons from $J = 2$ to $J = 5$. For every resulting size, we go through the same steps 1-3 as outlined above, and record the same statistics, displayed in Table 3 and Table 4, respectively. Essentially the same observations as before hold. Namely, policies of higher degree result in strict improvements of the objective function, with cubic policies always resulting in gaps smaller than 1% (see Figure 3(a) for a typical case). Also, increasing the problem size (here, this corresponds to the number of echelons, J) does not affect the reductions in gaps, and the computational requirements do not increase drastically (see Figure 3(b), which corresponds to quadratic policies).

All our computations were done in a MATLAB[®] environment, on the MIT Operations Research Center computational machine (3 GHz Intel[®] Dual Core Xeon[®] 5050 Processor, with 8GB of RAM

memory, running Ubuntu Linux). The optimization problems were formulated using YALMIP Löfberg [2004], and the resulting SDPs were solved with SDPT3 Toh et al. [1999].

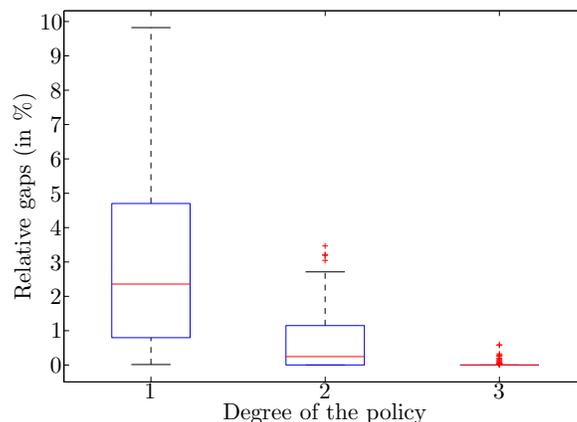
Table 1: Relative gaps (in %) for polynomial policies in Example 1

T	Degree $d = 1$					Degree $d = 2$					Degree $d = 3$				
	avg	std	mdn	min	max	avg	std	mdn	min	max	avg	std	mdn	min	max
4	2.84	2.41	2.18	0.02	9.76	0.75	0.85	0.47	0.00	3.79	0.03	0.12	0.00	0.00	0.91
5	2.82	2.29	2.52	0.04	11.22	0.62	0.71	0.39	0.00	3.92	0.02	0.09	0.00	0.00	0.56
6	3.09	2.63	2.36	0.01	9.82	0.69	0.89	0.25	0.00	3.47	0.03	0.10	0.00	0.00	0.59
7	3.25	2.95	2.58	0.13	15.00	0.83	0.99	0.43	0.00	4.79	0.06	0.17	0.00	0.00	0.93
8	3.66	3.29	2.69	0.03	18.36	1.06	1.17	0.74	0.00	5.81	0.10	0.17	0.00	0.00	0.99
9	2.93	2.78	2.12	0.05	11.56	0.80	0.86	0.55	0.00	3.39	0.07	0.13	0.00	0.00	0.61
10	3.44	3.60	2.09	0.00	18.20	0.76	1.16	0.26	0.00	5.76	0.05	0.12	0.00	0.00	0.74

Table 2: Solver times (in seconds) for polynomial policies in Example 1

T	Degree $d = 1$					Degree $d = 2$					Degree $d = 3$				
	avg	std	mdn	min	max	avg	std	mdn	min	max	avg	std	mdn	min	max
4	0.47	0.05	0.46	0.38	0.63	1.27	0.10	1.27	1.13	1.62	3.33	0.21	3.24	3.01	4.03
5	0.58	0.06	0.58	0.46	0.75	2.03	0.20	1.97	1.69	2.65	7.51	0.91	7.27	6.58	12.08
6	0.73	0.11	0.72	0.62	1.50	2.29	0.22	2.28	1.87	3.26	18.96	2.54	18.25	16.07	31.86
7	0.88	0.08	0.87	0.72	1.07	3.08	0.23	3.10	2.47	3.67	48.83	5.63	47.99	40.65	74.09
8	1.13	0.12	1.11	0.94	1.92	4.79	0.32	4.75	3.97	5.96	157.73	20.67	153.91	126.15	217.80
9	1.53	0.17	1.51	1.27	2.66	7.65	0.51	7.65	6.10	9.59	420.75	60.10	411.09	334.71	760.13
10	1.31	0.15	1.30	1.07	2.19	14.77	1.24	14.80	11.81	18.57	1846.94	600.89	1640.10	1313.18	4547.09

Figure 1: Box plots comparing the performance of different polynomial policies for horizon $T = 6$



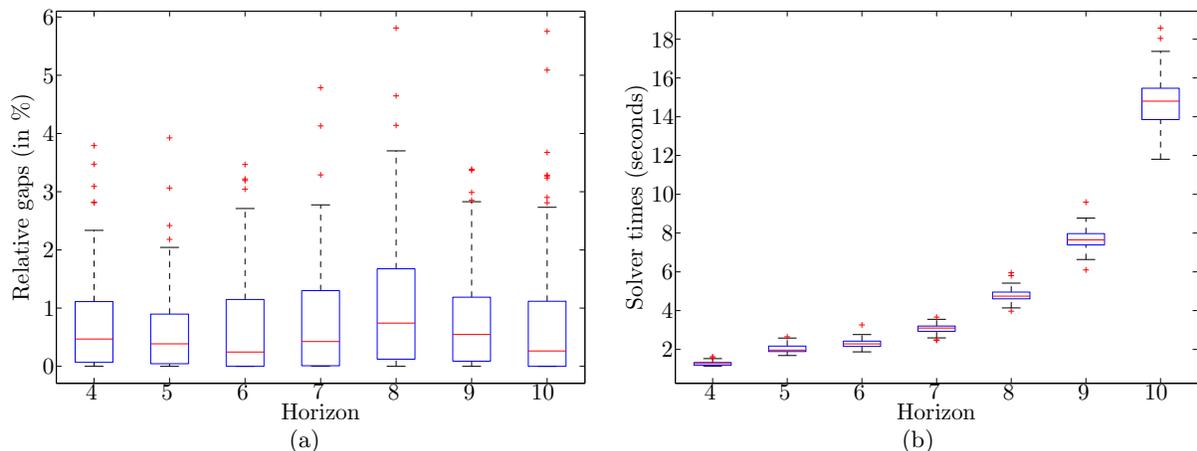


Figure 2: Performance of quadratic policies for Example 1 - (a) illustrates the weak dependency of the improvement on the problem size (measured in terms of the horizon T), while (b) compares the solver times required for different problem sizes.

Table 3: Relative gaps (in %) for polynomial policies in Example 2

J	Degree $d = 1$					Degree $d = 2$					Degree $d = 3$				
	avg	std	mdn	min	max	avg	std	mdn	min	max	avg	std	mdn	min	max
2	1.87	1.48	1.47	0.00	8.27	1.38	1.16	1.11	0.00	6.48	0.06	0.14	0.01	0.00	0.96
3	1.47	0.89	1.27	0.16	4.46	1.08	0.68	0.93	0.14	3.33	0.04	0.06	0.00	0.00	0.32
4	1.14	2.46	0.70	0.05	24.63	0.67	0.53	0.53	0.01	2.10	0.04	0.07	0.00	0.00	0.38
5	0.35	0.37	0.21	0.03	1.85	0.27	0.32	0.15	0.00	1.59	0.02	0.03	0.00	0.00	0.15

6 Conclusions

In this paper, we have presented a new method for dealing with multi-stage decision problems affected by uncertainty, applicable to robust optimization and stochastic programming. Our approach consists of constructing a hierarchy of sub-optimal polynomial policies, parameterized directly in the observed uncertainties. The problem of computing such an optimal polynomial policy can be reformulated as an SDP, which can be solved efficiently with interior point methods. Furthermore, the approach allows modelling flexibility, in that the degree of the polynomial policies explicitly controls the trade-off between the quality of the approximation and the computational requirements. To test the quality of the policies, we have considered two applications in inventory management, one involving a single echelon with constrained cumulative orders, and the second involving a serial supply chain. For both examples, quadratic policies (requiring modest computational requirements) were able to substantially reduce the optimality gap, and cubic policies (under more computational requirements) were always within 1% of optimal. Given that our tests were run using publicly-available, general-purpose SDP solvers, we believe that, with the advent of more powerful (commercial) packages for interior point methods, as well as dedicated algorithms for solving SOS problems, our method should have applicability to large scale, real-world optimization problems.

Table 4: Solver times (in seconds) for polynomial policies Example 2

J	Degree $d = 1$					Degree $d = 2$					Degree $d = 3$				
	avg	std	mdn	min	max	avg	std	mdn	min	max	avg	std	mdn	min	max
2	1.22	0.20	1.18	0.86	2.35	5.58	1.05	5.44	3.82	8.79	81.64	14.02	80.88	52.55	116.56
3	1.72	0.26	1.70	1.21	3.09	8.84	1.40	8.53	6.83	13.19	115.08	20.91	109.96	77.29	183.84
4	1.57	0.22	1.55	1.20	2.85	12.59	1.63	12.44	8.86	17.86	160.05	19.34	159.29	82.11	207.56
5	2.59	1.46	1.97	1.51	8.18	18.97	6.59	17.59	13.21	63.71	250.43	109.96	227.56	144.54	952.37

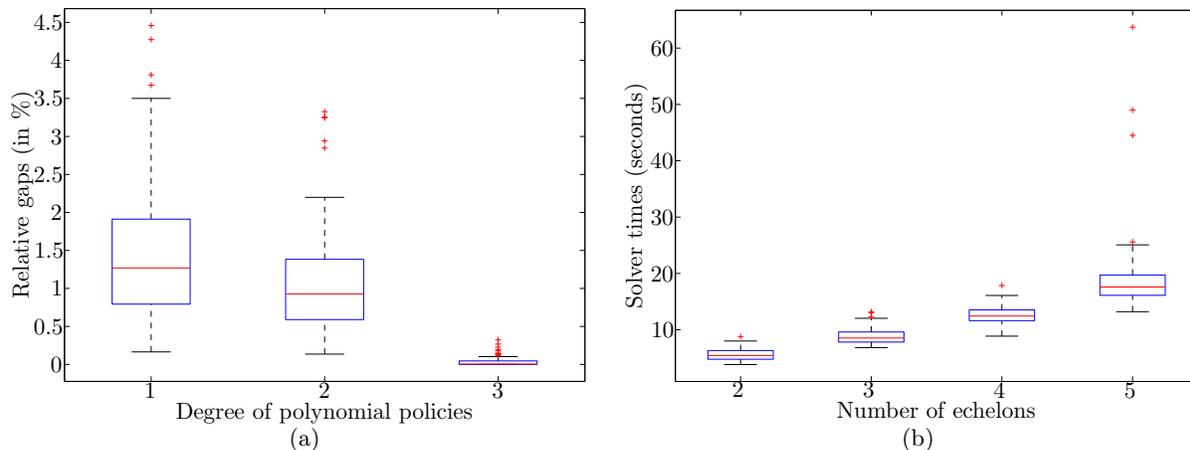


Figure 3: Performance of polynomial policies for Example 2. (a) compares the three policies for problems with $J = 3$ echelons, and (b) shows the solver times needed to compute quadratic policies for different problem sizes.

7 Appendix

7.1 Suboptimality of Affine Policies

Lemma 1. Consider Problem (9), written below for convenience. Recall that x is a (first-stage) non-adjustable decision, while \mathbf{u} is a second-stage adjustable policy (allowed to depend on \mathbf{w}).

$$\begin{aligned} & \underset{x, \mathbf{u}(\mathbf{w})}{\text{minimize}} && x \\ & \text{such that} && x \geq \sum_{i=1}^N u_i, \quad \forall \mathbf{w} \in \mathcal{W} = \{(w_1, \dots, w_N) \in \mathbb{R}^N : \|\mathbf{w}\|_2 \leq 1\}, \end{aligned} \quad (22a)$$

$$u_i \geq w_i^2, \quad \forall \mathbf{w} \in \mathcal{W}. \quad (22b)$$

The optimal value in the problem is 1, corresponding to policies $u_i(\mathbf{w}) = w_i^2$, $i = 1, \dots, N$. Furthermore, the optimal achievable objective under affine policies $\mathbf{u}(\mathbf{w})$ is N .

Proof. Note that for any feasible x, \mathbf{u} , we have $x \geq \sum_{i=1}^N u_i \geq \sum_{i=1}^N w_i^2$, for any $\mathbf{w} \in \mathcal{W}$. Therefore, with $\sum_{i=1}^N w_i^2 = 1$, we must have $x \geq 1$. Also note that $u_i^*(\mathbf{w}) = w_i^2$ is robustly feasible for constraint (22b), and results in an objective $x^* = \max_{\mathbf{w} \in \mathcal{W}} \sum_{i=1}^N w_i^2 = 1$, which equals the lower bound, and is hence optimal.

Consider an affine policy in the second stage, $u_i^{\text{AFF}}(\mathbf{w}) = \beta_i + \boldsymbol{\alpha}_i^T \mathbf{w}$, $i = 1, \dots, N$. With \mathbf{e}_1

denoting the first unit vector (1 in the first component, 0 otherwise), for any $i = 1, \dots, N$, we have:

$$\left. \begin{aligned} \mathbf{w} = \mathbf{e}_1 \in \mathcal{W} &\Rightarrow \beta_i + \alpha_i(1) \geq 1 \\ \mathbf{w} = -\mathbf{e}_1 \in \mathcal{W} &\Rightarrow \beta_i - \alpha_i(1) \geq 1 \end{aligned} \right\} \Rightarrow \beta_i \geq 1.$$

This implies that $x^{\text{AFF}} \geq \sum_{i=1}^N u_i^{\text{AFF}}(\mathbf{w}) \geq N + \sum_{i=1}^N \alpha_i^T \mathbf{w}$. In particular, with $\mathbf{w} = \mathbf{0} \in \mathcal{W}$, we have $x^{\text{AFF}} \geq N$. The optimal choice, in this case, will be to set $\alpha_i = \mathbf{0}$, resulting in $x^{\text{AFF}} = N$. \square

7.2 Optimality of Multi-affine Policies

Theorem 1. *Multi-affine policies of the form (19), with degree at most $d = T - 1$, are optimal for problem (P).*

Proof. The following trivial observation will be useful in our analysis:

Observation 1. *A multi-affine policy u_j of the form (19) is an affine function of a given variable ξ_i , when all the other variables ξ_l , $l \neq i$, are fixed. Also, with u_j of degree at most d , the number of coefficients l_α is $\binom{k}{0} + \binom{k}{1} + \dots + \binom{k}{d}$.*

Recall that the optimal value in Problem (P) is that same as the optimal value in Problem $(P)_{\text{ext}}$ from Section 3.3. Let us denote the optimal decisions obtained from solving problem $(P)_{\text{ext}}$ by $\mathbf{u}_k^{\text{ext}}(\boldsymbol{\xi})$, $\mathbf{x}_k^{\text{ext}}(\boldsymbol{\xi})$, respectively. Note that, at time k , there are at most 2^k such distinct values $\mathbf{u}_k^{\text{ext}}(\boldsymbol{\xi}_k)$, and, correspondingly, at most 2^k values $\mathbf{x}_k^{\text{ext}}(\boldsymbol{\xi}_k)$, due to the non-anticipativity condition and the fact that the extreme uncertainty sequences at time k , $\boldsymbol{\xi}_k \in \text{ext}(\Xi_k) = \text{ext}(\mathcal{W}_0) \times \dots \times \text{ext}(\mathcal{W}_{k-1})$, are simply the vertices of the hypercube $\Xi_k \subset \mathbb{R}^k$. In particular, at the last time when decisions are taken, $k = T - 1$, there are at most 2^{T-1} distinct optimal values $\mathbf{u}_{T-1}^{\text{ext}}(\boldsymbol{\xi}_{T-1})$ computed.

Consider now a multi-affine policy of the form (19), of degree $T - 1$, implemented at time $T - 1$. By Observation 1, the number of coefficients in the j -th component of such a policy is exactly $\binom{T-1}{0} + \binom{T-1}{1} + \dots + \binom{T-1}{T-1} = 2^{T-1}$, by Newton's binomial formula. Therefore, the total $n_u \cdot 2^{T-1}$ coefficients for \mathbf{u}_{T-1} could be computed so that

$$\mathbf{u}_{T-1}(\boldsymbol{\xi}_{T-1}) = \mathbf{u}_{T-1}^{\text{ext}}(\boldsymbol{\xi}_{T-1}), \quad \forall \boldsymbol{\xi}_{T-1} \in \text{ext}(\Xi_{T-1}), \quad (23)$$

that is, the value of the multi-affine policy exactly matches the 2^{T-1} optimal decisions computed in $(P)_{\text{ext}}$, at the 2^{T-1} vertices of Ξ_{T-1} . The same process can be conducted for times $k = T - 2, \dots, 1, 0$, to obtain multi-affine policies of degree at most¹⁰ $T - 1$ that match the values $\mathbf{u}_k^{\text{ext}}(\boldsymbol{\xi}_k)$ at the extreme points of Ξ_k .

With such multi-affine control policies, it is easy to see that the states \mathbf{x}_k become multi-affine functions of $\boldsymbol{\xi}_k$. Furthermore, we have $\mathbf{x}_k(\boldsymbol{\xi}_k) = \mathbf{x}_k^{\text{ext}}(\boldsymbol{\xi}_k)$, $\forall \boldsymbol{\xi}_k \in \text{ext}(\Xi_k)$. A typical state-control constraint (7c) written at time k amounts to ensuring that

$$\begin{aligned} \mathbf{e}_x(k, j)^T \mathbf{x}_k(\boldsymbol{\xi}_k) + \mathbf{e}_u(k, j)^T \mathbf{u}_k(\boldsymbol{\xi}_k) - \mathbf{f}_j(k) &\leq 0, \\ \forall \boldsymbol{\xi}_k \in \Xi_k = \mathcal{W}_0 \times \dots \times \mathcal{W}_{k-1}, \end{aligned}$$

where $\mathbf{e}_x(k, j)^T$, $\mathbf{e}_u(k, j)^T$ denote the j -th row of $E_x(k)$ and $E_u(k)$, respectively. Note that the left-hand side of this expression is also a multi-affine function of the variables $\boldsymbol{\xi}_k$. Since, by our observation, the maximum of multi-affine functions is reached at the vertices of the feasible set, i.e., $\boldsymbol{\xi}_k \in \text{ext}(\Xi_k)$, and, by (23), we have that for any such vertex, $u_k(\boldsymbol{\xi}_k) = \mathbf{u}_k^{\text{ext}}(\boldsymbol{\xi}_k)$, $\mathbf{x}_k(\boldsymbol{\xi}_k) = \mathbf{x}_k^{\text{ext}}(\boldsymbol{\xi}_k)$,

¹⁰In fact, multi-affine policies of degree k would be sufficient at time k

we immediately conclude that the constraint above is satisfied, since $\mathbf{u}_k^{\text{ext}}(\boldsymbol{\xi}_k), \mathbf{x}_k^{\text{ext}}(\boldsymbol{\xi}_k)$ are certainly feasible.

A similar argument can be invoked for constraint (7d), and also to show that the maximum of the objective function is reached on the set of vertices $\text{ext}(\Xi_T)$, and, since the values of the multi-affine policies exactly correspond to the optimal decisions in program $(P)_{\text{ext}}$, optimality is preserved. \square

References

- Soren Asmussen and Peter Glynn. *Stochastic Simulation: Algorithms and Analysis*. Springer, 2007. ISBN 978-0387306797.
- A. Bemporad, F. Borrelli, and M. Morari. Optimal Controllers for Hybrid Systems: Stability and Piecewise Linear Explicit Form. In *Decision and Control, 2000. Proceedings of the 39th IEEE Conference on*, volume 2, pages 1810–1815 vol.2, 2000. doi: 10.1109/CDC.2000.912125.
- A. Bemporad, F. Borrelli, and M. Morari. Model Predictive Control Based on Linear Programming - The Explicit Solution. *IEEE Transactions on Automatic Control*, 47(12):1974–1985, 2002.
- A. Bemporad, F. Borrelli, and M. Morari. Min-Max Control of Constrained Uncertain Discrete-Time Linear Systems. *IEEE Transactions on Automatic Control*, 48(9):1600–1606, 2003. ISSN 0018-9286.
- A. Ben-Tal and A. Nemirovski. Robust Optimization - Methodology and Applications. *Mathematical Programming*, 92(3):453–480, 2002. ISSN 0025-5610. doi: <http://dx.doi.org/10.1007/s101070100286>.
- A. Ben-Tal, A. Goryashko, E. Guslitzer, and A. Nemirovski. Adjustable Robust Solutions of Uncertain Linear Programs. *Mathematical Programming*, 99(2):351–376, 2004. ISSN 0025-5610. doi: <http://dx.doi.org/10.1007/s10107-003-0454-y>.
- A. Ben-Tal, S. Boyd, and A. Nemirovski. Control of Uncertainty-Affected Discrete Time Linear Systems via Convex Programming. Working paper, 2005a.
- A. Ben-Tal, B. Golany, A. Nemirovski, and J.-P. Vial. Retailer-Supplier Flexible Commitments Contracts: A Robust Optimization Approach. *Manufacturing & Service Operations Management*, 7(3):248–271, 2005b. ISSN 1526-5498. doi: <http://dx.doi.org/10.1287/msom.1050.0081>.
- A. Ben-Tal, S. Boyd, and A. Nemirovski. Extending Scope of Robust Optimization: Comprehensive Robust Counterparts of Uncertain Problems. *Mathematical Programming*, 107(1):63–89, 2006. ISSN 0025-5610. doi: <http://dx.doi.org/10.1007/s10107-005-0679-z>.
- D. P. Bertsekas. *Dynamic Programming and Optimal Control*. Athena Scientific, Belmont, MA, 2001. ISBN 1886529272.
- D. Bertsimas and D.B. Brown. Constrained Stochastic LQC: A Tractable Approach. *IEEE Transactions on Automatic Control*, 52(10):1826–1841, Oct. 2007. ISSN 0018-9286. doi: <http://dx.doi.org/10.1109/TAC.2007.906182>.
- D. Bertsimas, D. Iancu, and P. Parrilo. Optimality of Affine Policies in Robust Optimization. Submitted for publication, 2009. URL <http://arxiv.org/abs/0904.3986>.
- John R. Birge and Francois Louveaux. *Introduction to Stochastic Programming*. Springer, 2000. ISBN 0387982175.
- Debasish Chatterjee, Peter Hokayem, and John Lygeros. Stochastic Model Predictive Control with Bounded Control Inputs: A Vector Space Approach. Submitted for publication., March 2009. URL <http://arxiv.org/abs/0903.5444>.

- Xin Chen and Yuhan Zhang. Uncertain Linear Programs: Extended Affinely Adjustable Robust Counterparts. *Operations Research*, page opre.1080.0605, 2009. doi: 10.1287/opre.1080.0605. URL <http://or.journal.informs.org/cgi/content/abstract/opre.1080.0605v1>.
- Xin Chen, Melvyn Sim, Peng Sun, and Jiawei Zhang. A Linear Decision-Based Approximation Approach to Stochastic Programming. *Operations Research*, 56(2):344–357, 2008. doi: 10.1287/opre.1070.0457. URL <http://or.journal.informs.org/cgi/content/abstract/56/2/344>.
- D. P. de Farias and B. Van Roy. The Linear Programming Approach to Approximate Dynamic Programming. *Operations Research*, 51(6):850–865, 2003. doi: 10.1287/opre.51.6.850.24925. URL <http://or.journal.informs.org/cgi/content/abstract/51/6/850>.
- Daniela Pucci de Farias and Benjamin Van Roy. On Constraint Sampling in the Linear Programming Approach to Approximate Dynamic Programming. *Mathematics of Operations Research*, 29(3):462–478, 2004. doi: 10.1287/moor.1040.0094. URL <http://mor.journal.informs.org/cgi/content/abstract/29/3/462>.
- G. E. Dullerud and F. Paganini. *A Course in Robust Control Theory*. Springer, 2005. ISBN 0387989455.
- Martin Dyer and Leen Stougie. Computational Complexity of Stochastic Programming Problems. *Mathematical Programming*, 106(3):423–432, 2006. ISSN 0025-5610. doi: <http://dx.doi.org/10.1007/s10107-005-0597-0>.
- Stanley J. Gartska and Roger J.-B. Wets. On Decision Rules in Stochastic Programming. *Mathematical Programming*, 7(1):117–143, 1974. ISSN 0025-5610. doi: 10.1007/BF01585511.
- P. J. Goulart and E.C. Kerrigan. Relationships Between Affine Feedback Policies for Robust Control with Constraints. In *Proceedings of the 16th IFAC World Congress on Automatic Control*, July 2005.
- T. Jacobi and A. Prestel. Distinguished Representations of Strictly Positive Polynomials. *Journal für die reine und angewandte Mathematik (Crelles Journal)*, 2001(532):223–235, 2001. doi: 10.1515/crll.2001.023. URL <http://www.reference-global.com/doi/abs/10.1515/crll.2001.023>.
- E.C. Kerrigan and J.M. Maciejowski. On Robust Optimization and the Optimal Control of Constrained Linear Systems with Bounded State Disturbances. In *Proceedings of the 2003 European Control Conference*, Cambridge, UK, September 2003.
- E.C. Kerrigan and J.M. Maciejowski. Properties of a New Parameterization for the Control of Constrained Systems with Disturbances. *Proceedings of the 2004 American Control Conference*, 5:4669–4674 vol.5, June-July 2004. ISSN 0743-1619. URL <http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=01384049>.
- J. B. Lasserre. Global Optimization with Polynomials and the Problem of Moments. *SIAM Journal on Optimization*, 11:796–817, 2001.
- J. Löfberg. Approximations of Closed-loop Minimax MPC. *Proceedings of the 42nd IEEE Conference on Decision and Control*, 2:1438–1442 Vol.2, December 2003. ISSN 0191-2216. doi: 10.1109/CDC.2003.1272813.
- J. Löfberg. YALMIP : A Toolbox for Modeling and Optimization in MATLAB. In *Proceedings of the CACSD Conference*, Taipei, Taiwan, 2004. URL <http://control.ee.ethz.ch/~joloef/yalmip.php>.
- Peter Marbach and John N. Tsitsiklis. Simulation-based Optimization of Markov Reward Processes. *IEEE Transactions on Automatic Control*, 46(2):191–209, January 2001. doi: 10.1109/9.905687.
- Arkadi Nemirovski and Alexander Shapiro. *Continuous Optimization: Current Trends and Applications*, chapter On Complexity of Stochastic Programming Problems, pages 111–146. Springer US, 2005. doi: 10.1007/b137941.

- Pablo A. Parrilo. *Structured Semidefinite Programs and Semialgebraic Geometry Methods in Robustness and Optimization*. PhD thesis, California Institute of Technology, 2000.
- Pablo A. Parrilo. Semidefinite Programming Relaxations for Semialgebraic Problems. *Mathematical Programming Series B*, 96(2):293–320, May 2003. doi: <http://dx.doi.org/10.1007/s10107-003-0387-5>.
- M. Putinar. Positive Polynomials on Compact Semi-algebraic Sets. *Indiana University Mathematics Journal*, 42(3):969–984, 2003.
- Walter Rudin. *Principles of Mathematical Analysis*. Mc-Graw Hill Science/Engineering/Math, 1976. ISBN 007054235X.
- J. Skaf and S. Boyd. Design of Affine Controllers Via Convex Optimization. Submitted to *IEEE Transactions on Automatic Control*, 2008a. URL http://www.stanford.edu/~boyd/papers/affine_contr.html.
- J. Skaf and S. Boyd. Nonlinear Q-Design for Convex Stochastic Control. Submitted to *IEEE Transactions on Automatic Control*, 2008b. URL http://www.stanford.edu/~boyd/papers/nonlin_Q_param.html.
- K.C. Toh, M.J. Todd, and R. Tütüncü. SDPT3 a MATLAB Software Package for Semidefinite Programming, 1999. URL <http://www.math.nus.edu.sg/~mattohkc/sdpt3.html>.
- Lieven Vandenberghe and Stephen Boyd. Semidefinite Programming. *SIAM Review*, 38(1):49–95, 1996. ISSN 0036-1445. doi: <http://dx.doi.org/10.1137/1038003>.
- K. Zhou and J. C. Doyle. *Essentials of Robust Control*. Prentice Hall, 1998. ISBN 0135258332.
- P. Zipkin. *Foundations of Inventory Management*. McGraw Hill, 2000. ISBN 0256113793.