The Value of Adaptability

Dimitris Bertsimas * Constantine Caramanis†

September 30, 2005

Abstract

We consider linear optimization problems with deterministic parameter uncertainty. We consider a departure from the robust optimization paradigm, allowing the decision-maker some limited, finite adaptability. Equivalently, we can view this as though the decision-maker can obtain some additional information about the uncertainty before committing to a decision. The central problem we address is optimally structuring finite adaptability, and understanding the marginal value of adaptability. We propose a hierarchy of increasing adaptability that bridges the gap between the robust formulation, and the re-optimization (or fully adjustable) formulation. We show some negative complexity results; then we provide an efficient (heuristic) algorithm for adaptability, and provide some computational evidence of its efficacy. The framework we propose is also appropriate for problems with discrete variables. To the best of our knowledge, this is the first proposal for adaptability that can accommodate discrete variables. We also compare our proposal to the model of affine adaptability proposed in [2].

1 Introduction

Recently, much work has been done in robust optimization, in the case of linear, semidefinite, and general conic programming, and also for discrete combinatorial robust optimization; see, e.g., [3],[4], [6], [7], [11]. The landscape of solution concepts has two extreme cases. On the one side, we have the robust formulation where the decision maker has no adaptability to, or information about, the realization of the uncertainty. In some settings this may be overly conservative, particularly when uncertainty affecting different constraints is not independent. On the other extreme is the re-optimization formulation where the decision-maker has advance knowledge of the exact realization of the uncertainty (or equivalently, complete adaptability) and then selects an optimal solution accordingly. This set-up is overly optimistic. Full information is rarely available, and exact observations rarely possible. Moreover, even if in principle feasible, implementing complete adaptability may be too expensive, and hence in itself undesirable. This motivates us to consider the middle ground.

Consider the case where the decision-maker has the ability (at some cost) to obtain some coarse information about the realization of the uncertainty, or (again at some cost) is able to select some finite number, k, of contingency plans (as opposed to a single robust solution). The central topic of this paper is the value of this adaptability.

Contributions and Paper Outline

In Section 2 we provide the basic setup of our adaptability proposal. Because of its inherent

*Boeing Professor of Operations Research, Sloan School of Management and Operations Research Center, Massachusetts Institute of Technology, E53-363, Cambridge, MA 02142, dbertsim@mit.edu

†Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA 02139-4307; cmcaram@mit.edu. Corresponding author.

discrete nature, this adaptability proposal can accommodate discrete variables. To the best of our knowledge, this is the first proposal for adaptability that can reasonably deal with discrete variables. In Section 3 we give a geometric interpretation of the conservativeness of the robust formulation. We provide a geometric characterization of when finite adaptability can improve the robust solution by η, for any (possibly large) chosen $\eta \geq 0$. We obtain (collectively sufficient) necessary conditions that any k-adaptability scheme must satisfy in order to improve the robust solution by at least η.

In Section 4 we also consider affine adaptability. Following the work of Ben-Tal et al. in [2], there has been renewed interest in adaptability (e.g., [1],[9],[14]). Our work differs from affine adaptability proposals in several important ways. First, the model offers a natural hierarchy of increasing adaptability. Second, the intrinsic discrete aspect of the adaptability proposal makes this suitable for any situation where it may not make sense to be sensitive to infinitesimal changes in the data. Indeed, only coarse observations may be available. In addition, especially from a control viewpoint, infinite (and thus infinitesimal) adjustability as required by the affine adaptability framework, may not be feasible, or even desirable. We provide an example where affine adaptability is no better than the robust solution, while finite adaptability with 3 contingency plans significantly improves the solution.

In Section 5 we consider the special case of right hand side uncertainty. We provide an IP formulation for the $k = 2$ contingency plan problem. We also show here that structuring $k = 2$ adaptability optimally, is NP-hard. In Section 6 we provide an approximation algorithm based on the qualitative prescriptions of Section 3. Section 7 provides several computational examples.

2 The Basic Setup

We consider linear programming problems with deterministic uncertainty in the coefficients, where the uncertainty set is polyhedral. Uncertainty in the right hand side or in the objective function can be modeled by uncertainty in the matrix (see, e.g., [6]). Thus the general problem we consider is

$$\text{Robust}(\mathcal{P}) \triangleq \left[\begin{array}{ll} \min: & \boldsymbol{c}'\boldsymbol{x} \\ \text{s.t.}: & \boldsymbol{A}\boldsymbol{x} \geq \boldsymbol{b}, \quad \forall \boldsymbol{A} \in \mathcal{P} \end{array} \right], \tag{2.1}$$

where there are m constraints, n variables, and $\mathcal{P}$ is a polytope. We consider both the case where $\mathcal{P}$ is given as a convex hull of its extreme points, and where it is given as the intersection of half-spaces. Some results apply only to the case of the convex hull representation; we make this clear when it is the case. The re-optimization formulation is

$$\text{ReOpt}(\mathcal{P}) \triangleq \max_{\boldsymbol{A} \in \mathcal{P}} \left[\begin{array}{ll} \min: & \boldsymbol{c}'\boldsymbol{x} \\ \text{s.t.}: & \boldsymbol{A}\boldsymbol{x} \geq \boldsymbol{b} \end{array} \right].$$

In the k-adaptability problem, the decision-maker chooses k solutions, $\{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_k\}$, and then commits to one of them only after seeing the realization of the uncertainty set. Following an algebraic formulation, k-adaptability is given by:

$$\text{Adapt}_k(\mathcal{P}) \triangleq \left[\begin{array}{ll} \min: & \max\{\boldsymbol{c}'\boldsymbol{x}_1, \ldots, \boldsymbol{c}'\boldsymbol{x}_k\} \\ \text{s.t.}: & [\boldsymbol{A}\boldsymbol{x}_1 \geq \boldsymbol{b} \ \text{ or } \ \boldsymbol{A}\boldsymbol{x}_2 \geq \boldsymbol{b} \ \text{ or } \ \cdots \ \text{ or } \ \boldsymbol{A}\boldsymbol{x}_k \geq \boldsymbol{b}] \quad \forall \boldsymbol{A} \in \mathcal{P} \end{array} \right].$$

This is a disjunctive optimization problem with infinitely many constraints. We formulate this as a (finite) bilinear optimization problem. For notational convenience, we consider the case $k = 2$, but the extension to the general case is straightforward. For this reformulation, we focus on the case where the uncertainty set $\mathcal{P}$ is given as a convex hull of its extreme points: $\mathcal{P} = \text{conv}\{\boldsymbol{A}^1, \ldots, \boldsymbol{A}^K\}$.

Proposition 1 *The optimal 2-adaptability value, and the optimal two contingency plans are given by the solution to the bilinear optimization:*

$$\begin{array}{rl} \min: & \boldsymbol{c}'\boldsymbol{x}_1 \vee \boldsymbol{c}'\boldsymbol{x}_2 \\ \text{s.t.}: & \mu_{ij}\left[(\boldsymbol{A}^l\boldsymbol{x}_1)_i - b_i\right] + (1-\mu_{ij})\left[(\boldsymbol{A}^l\boldsymbol{x}_2)_j - b_j\right] \geq 0, \quad \forall 1 \leq i,j \leq m, \quad \forall 1 \leq l \leq K \\ & 0 \leq \mu_{ij} \leq 1, \quad \forall 1 \leq i,j \leq m. \end{array}$$

The proof is in the full version, [5]. Bilinear optimization is typically hard, but there has been much work towards algorithmic solutions (see, e.g., [10],[13], and references therein). Equivalently we can follow a geometric formulation. Consider, rather than having some limited adaptability, that the decision-maker receives some advance side-information about the realization of the uncertainty. The decision-maker selects a partition of the uncertainty set $\mathcal{P}$ into k (possibly non-disjoint) regions: $\mathcal{P} = \mathcal{P}_1 \cup \cdots \cup \mathcal{P}_k$. Then before committing to a single solution, the decision-maker receives information about which of the k regions of the partition contains the realization of the uncertainty. The optimal k-adaptability problem becomes

$$\text{Adapt}_k(\mathcal{P}) \stackrel{\triangle}{=} \min_{\mathcal{P}=\mathcal{P}_1\cup\cdots\cup\mathcal{P}_k} \left[\begin{array}{rll} \min: & \max\{\boldsymbol{c}'\boldsymbol{x}_1, \ldots, \boldsymbol{c}'\boldsymbol{x}_k\} & \\ \text{s.t.}: & \boldsymbol{A}\boldsymbol{x}_1 \geq b, & \forall \boldsymbol{A} \in \mathcal{P}_1 \\ & \vdots & \\ & \boldsymbol{A}\boldsymbol{x}_k \geq b, & \forall \boldsymbol{A} \in \mathcal{P}_k. \end{array} \right].$$

Throughout this paper we refer equivalently to either k contingency plans, or k-partitions, for the k-adaptability problem. The inequalities $\text{Robust}(\mathcal{P}) \geq \text{Adapt}_k(\mathcal{P}) \geq \text{ReOpt}(\mathcal{P})$ hold in general.

3 A Geometric View and Necessary Conditions

In this section we provide a geometric view of the gap between the re-optimization and robust formulations, and also of the way in which finite adaptability bridges this gap. Then, we use this geometric interpretation to obtain necessary conditions that any k-partition must satisfy in order to improve the robust solution value by at least η, for any chosen value η.

3.1 The Geometric Gap

Given any uncertainty region $\mathcal{P}$, let $\pi_l(\mathcal{P})$ denote the projection of $\mathcal{P}$ onto the components corresponding to the l^{th} constraint of (2.1), $1 \leq l \leq m$. Then, define:

$$(\mathcal{P})_R \stackrel{\triangle}{=} \text{conv}(\pi_1\mathcal{P}) \times \text{conv}(\pi_2\mathcal{P}) \times \cdots \times \text{conv}(\pi_m\mathcal{P}). \tag{3.2}$$

The set $(\mathcal{P})_R$ is the smallest hypercube (in the above sense) that contains the set $\mathcal{P}$.

Lemma 1 *For $\mathcal{P}$, and $(\mathcal{P})_R$ defined as above, we have:*
(1) $\text{Robust}(\mathcal{P}) = \text{Robust}((\mathcal{P})_R)$ *and* $\text{Robust}(\mathcal{P}) = \text{ReOpt}((\mathcal{P})_R)$.
(2) For $\mathcal{P} = \mathcal{P}_1 \cup \cdots \cup \mathcal{P}_k$ the optimal k-partition of the uncertainty set,

$$\text{Adapt}_k(\mathcal{P}) = \text{ReOpt}((\mathcal{P}_1)_R \cup \cdots \cup (\mathcal{P}_k)_R).$$

(3) There is a sequence of partitions $\{\mathcal{P}_{k1} \cup \mathcal{P}_{k2} \cup \cdots \cup \mathcal{P}_{kk}\}$ so that

$$(\mathcal{P}_{k1})_R \cup \cdots \cup (\mathcal{P}_{kk})_R \longrightarrow \mathcal{P}, \quad k \to \infty.$$

The first part of the lemma says that the robust formulation cannot model nonconvexity, or correlation across different constraints, and this accounts for the gap between the robust and the re-optimization formulations. The second and third parts of the lemma explain from a geometric perspective why, and how, the adaptive solution improves the robust cost.

PROOF. Assertions (1) and (2) follow from LP duality ([4]). For (3), consider any sequence of partitions where the maximum diameter of any region goes to zero as $k \to \infty$. The full details are in [5]. $\square$

We would like to conclude that k-adaptability bridges the gap between the robust and re-optimization values, i.e., $\text{Adapt}_k(\mathcal{P}) \to \text{ReOpt}(\mathcal{P})$ as k increases. With an additional continuity assumption, the proposition below asserts that this is in fact the case.

Continuity Assumption: *For any $\varepsilon > 0$, for any $\boldsymbol{A} \in \mathcal{P}$, there exists $\delta > 0$ and a point $\boldsymbol{x}$, feasible for $\boldsymbol{A}$ and within ε of optimality, such that $\forall\ \boldsymbol{A}' \in \mathcal{P}$ with $d(\boldsymbol{A}, \boldsymbol{A}') \leq \delta$, $\boldsymbol{x}$ is also feasible for $\boldsymbol{A}'$.*

The Continuity Assumption is fairly mild. It asks that the set of points feasible for two infinitesimally close matrices $\boldsymbol{A}, \boldsymbol{A}'$, (here $d(\cdot,\cdot)$ is the usual metric) be relatively large.

Proposition 2 *If our problem satisfies the Continuity Assumption, then given any sequence of partitions $\{((\mathcal{P}_{k1})_R \cup \cdots \cup (\mathcal{P}_{kk})_R)\}_{k=1}^{\infty}$ with the diameter of the largest set going to zero,*

$$\lim_{k\to\infty} \text{Adapt}_k(\mathcal{P}) = \text{ReOpt}(\mathcal{P}).$$

PROOF. Given any $\varepsilon > 0$, for every $\boldsymbol{A} \in \mathcal{P}$, consider the $\delta(\boldsymbol{A})$-neighborhood around $\boldsymbol{A}$ as given by the Continuity Assumption. These neighborhoods form an open cover of $\mathcal{P}$. Since $\mathcal{P}$ is compact, we can select a finite subcover. Let the partition $\mathcal{P} = \mathcal{P}_1 \cup \cdots \cup \mathcal{P}_k$ be (the closure of) such a subcover. Then, by the Continuity Assumption, $\text{Robust}(\mathcal{P}_i) \leq \text{ReOpt}(\mathcal{P}_i) + \varepsilon$. By definition, $\text{ReOpt}(\mathcal{P}_i) \leq \max_j \text{ReOpt}(\mathcal{P}_j) = \text{ReOpt}(\mathcal{P})$. From here the result follows. In [5] we show that the Continuity Assumption cannot be removed. $\square$

3.2 Necessary Conditions for η-Improvement

Lemma 1 says that $\text{Robust}(\mathcal{P}) = \text{ReOpt}((\mathcal{P})_R)$. Therefore, there must exist some $\hat{\boldsymbol{A}} \in (\mathcal{P})_R$ for which the nominal LP has value equal to the robust optimal value. Let $\mathcal{A}$ denote all such matrices. In fact, we show that for any $\eta > 0$, there exists a set $\mathcal{A}_\eta \subseteq (\mathcal{P})_R$ such that if $\boldsymbol{c}'\boldsymbol{x} \leq \text{Robust}(\mathcal{P}) - \eta$, then $\boldsymbol{x}$ does not satisfy $\boldsymbol{A}\boldsymbol{x} \geq \boldsymbol{b}$, for any $\boldsymbol{A} \in \mathcal{A}_\eta$.

Proposition 3 *(1) The sets $\mathcal{A}$ and $\mathcal{A}_\eta$ are the images under a computable map, of a polytope associated with the LP dual of the robust problem.*

(2) Adaptability with k contingency plans corresponding to the partition $\mathcal{P} = \mathcal{P}_1 \cup \cdots \cup \mathcal{P}_k$ improves the cost by at least η if and only if $((\mathcal{P}_1)_R \cup \cdots \cup (\mathcal{P}_k)_R) \cap \bar{\mathcal{A}}_\eta = \emptyset$.

(3) There is some k for which optimally chosen contingency plans can improve the cost by at least η if and only if $\bar{\mathcal{A}}_\eta \cap \mathcal{P} = \emptyset$.

The proof depends on Lemma 2 below. The details are postponed to [5],[5]. For the remainder of this section, we consider the case where the uncertainty is given as the convex hull of a given set of extreme points: $\mathcal{P} = \text{conv}\{\boldsymbol{A}^1, \ldots, \boldsymbol{A}^K\}$. For any $\eta > 0$, we consider the infeasible problem

$$\begin{array}{rl} \min: & \mathbf{0} \\ \text{s.t.}: & \boldsymbol{A}^i\boldsymbol{x} \geq \boldsymbol{b}, \quad 1 \leq i \leq K \\ & \boldsymbol{c}'\boldsymbol{x} \leq \text{Robust}(\mathcal{P}) - \eta. \end{array}$$

The dual is feasible, and hence unbounded. Let $\mathcal{C}_\eta$ be the set of directions of dual unboundedness, and $\mathcal{C}_0$ the set of dual optimal solutions:

$$\mathcal{C}_\eta \triangleq \left\{ (\boldsymbol{p}_1, \ldots, \boldsymbol{p}_K) : \begin{array}{l} (\boldsymbol{p}_1 + \cdots + \boldsymbol{p}_K)'\boldsymbol{b} \geq \text{Robust}(\mathcal{P}) - \eta \\ \boldsymbol{p}_1'\boldsymbol{A}^1 + \cdots + \boldsymbol{p}_K'\boldsymbol{A}^K = \boldsymbol{c} \\ \boldsymbol{p}_1, \ldots, \boldsymbol{p}_K \geq 0. \end{array} \right\}$$

For $(\boldsymbol{p}_1, \ldots, \boldsymbol{p}_K) \in \mathcal{C}_\eta$ let p_{ij} denote the j^{th} component of $\boldsymbol{p}_i$. Let $(\boldsymbol{A}^i)_j$ denote the j^{th} row of the matrix $\boldsymbol{A}^i$. Construct a matrix $\tilde{\boldsymbol{A}}$ whose j^{th} row is given by

$$(\tilde{\boldsymbol{A}})_j = \begin{cases} \boldsymbol{0} & \text{if } \sum_i p_{ij} = 0. \\ \frac{p_{1j}(\boldsymbol{A}^1)_j + \cdots + p_{Kj}(\boldsymbol{A}^K)_j}{\sum_i p_{ij}} & \text{otherwise.} \end{cases}$$

Thus each nonzero row of $\tilde{\boldsymbol{A}}$ is a convex combination of the corresponding rows of the $\boldsymbol{A}^i$ matrices. Let $\hat{\boldsymbol{A}}$ consist of any matrix in $(\mathcal{P})_R$ that coincides with $\tilde{\boldsymbol{A}}$ on all its non-zero rows. The proof of the next lemma follows by a straightforward duality argument.

Lemma 2 *For $\hat{\boldsymbol{A}}$ defined as above, the value of the optimal solution to the nominal LP is no less than* $\text{Robust}(\mathcal{P}) - \eta$. *If $\eta = 0$, and if $\boldsymbol{x}_R$ is an optimal solution for the robust problem, then $\boldsymbol{x}_R$ is also an optimal solution for the nominal problem with the matrix $\hat{\boldsymbol{A}}$.*

The collection of such $\hat{\boldsymbol{A}}$ make up $\mathcal{A}$ and $\mathcal{A}_\eta$. Thus $\mathcal{A}$ and $\mathcal{A}_\eta$ are images of $\mathcal{C}_0$ and $\mathcal{C}_\eta$, respectively, under the mapping given above. This concludes the proof of Proposition 3. More details about the sets $\mathcal{A}$, $\mathcal{A}_\eta$, and the mapping from $\mathcal{C}_\eta$ are contained in [5].

We now use the characterization of Proposition 3 to obtain necessary conditions that any η-improving partition must satisfy.

Proposition 4 *(1) Consider any element $\tilde{\boldsymbol{A}}$ obtained from a point of $\mathcal{C}_\eta$. Let us assume that the first r rows of the matrix $\tilde{\boldsymbol{A}}$ are nonzero. Let $\mathcal{Q}_j = \pi_j^{-1}(\tilde{\boldsymbol{A}}_j)$ denote the set of matrices in $\mathcal{P}$ whose j^{th} row equals the j^{th} row of $\tilde{\boldsymbol{A}}$. Then a partition $\mathcal{P} = \mathcal{P}_1 \cup \cdots \cup \mathcal{P}_k$ can achieve an improvement of at least η only if for any region $\mathcal{P}_d$, $1 \leq d \leq k$, there exists some $1 \leq i \leq r$, such that*

$$\mathcal{P}_d \cap \mathcal{Q}_i = \emptyset.$$

(2) Collectively, these necessary conditions are also sufficient.

PROOF. If there exists a region $\mathcal{P}_d$ such that $\mathcal{P}_d \cap \mathcal{Q}_i \neq \emptyset$ for $1 \leq i \leq r$, then we can find matrices $\boldsymbol{Q}_1, \ldots, \boldsymbol{Q}_r$, such that $\boldsymbol{Q}_i \in \mathcal{P}_d \cap \mathcal{Q}_i$. By definition, the i^{th} row of matrix $\boldsymbol{Q}_i$ coincides with the i^{th} row of $\tilde{\boldsymbol{A}}$. Therefore $\hat{\boldsymbol{A}} \in (\mathcal{P}_d)_R$. Now the proof of necessity follows from Proposition 3.

For sufficiency, failure to achieve the guaranteed improvement corresponds to a direction of dual unboundedness in one of the restricted problems, which in turn translates into a necessary condition for the original problem. The full details are in [5]. ☐

Therefore we can map any point of $\mathcal{C}_\eta$ to a necessary condition that any η-improving partition must satisfy. In Section 5 we show that computing the optimal partition is NP-hard. Nevertheless, necessary conditions allow the possibility of providing a short certificate that there does not exist a partition with $k \leq k'$, that achieves η-improvement. We provide a simple example of this phenomenon in the next section. Indeed, in this example, a finite (and small) set of necessary conditions reveals the limits, and structure of 2,3,4,5-adaptability.

4 An Extended Example

This section serves a dual purpose. We present a small example, illustrating the use of necessary conditions. We also illustrate that affine adaptability, even though it requires infinite adjustability, may do poorly even where 3-adaptability does quite well. We consider an example with one-dimensional uncertainty set.

$$\begin{aligned} \min: \quad & x_1 + x_2 + x_3 \\ \text{s.t.}: \quad & \boldsymbol{A}\boldsymbol{x} \geq \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} \qquad \forall\, \boldsymbol{A} \in \text{conv}\{\boldsymbol{A}^1, \boldsymbol{A}^2\} = \text{conv}\left\{ \begin{pmatrix} 0 & 0 & 1 \\ 1 & 1 & 1 \\ \frac{1}{2} & \frac{1}{5} & 0 \end{pmatrix}, \begin{pmatrix} 1 & 1 & 1 \\ 0 & 0 & 1 \\ \frac{1}{5} & \frac{1}{2} & 0 \end{pmatrix} \right\} \\ & x_1, x_2, x_3 \geq 0. \end{aligned}$$

The unique optimal solution is $\boldsymbol{x}_R = (10/7, 10/7, 1)$, with corresponding value $\text{Robust}(\mathcal{P}) = 27/7$. The re-optimization value is $\text{ReOpt}(\mathcal{P}) = 3$. We generate the necessary conditions corresponding to extreme points of the polytope $\mathcal{C}_0$, and also $\mathcal{C}_{\eta_1}, \mathcal{C}_{\eta_2}$, for $\eta_1 = (27/7) - 3.2$, and $\eta_2 = (27/7) - 2.9$. Following Proposition 4 we find points of $\mathcal{P}$ that any η-improving partition must separate. The two extreme points of $\mathcal{C}_0$ imply the two necessary conditions: $\mathcal{N}_1 = \{\boldsymbol{A}^2, \frac{1}{2}\boldsymbol{A}^1 + \frac{1}{2}\boldsymbol{A}^2\}$, and $\mathcal{N}_2 = \{\boldsymbol{A}^1, \frac{1}{2}\boldsymbol{A}^1 + \frac{1}{2}\boldsymbol{A}^2\}$. Evidently, no partition into 2 (convex) regions exists that satisfies both $\mathcal{N}_1$ and $\mathcal{N}_2$. Therefore 2-adaptability (in this example) is no better than the robust solution. The polytope $\mathcal{C}_{\eta_1}$ has 12 extreme points[1] which yield four non-redundant necessary conditions:

$$\begin{array}{ll} \mathcal{N}_1 = \{\boldsymbol{A}^1, \frac{28}{33}\boldsymbol{A}^1 + \frac{5}{33}\boldsymbol{A}^2\}; & \mathcal{N}_2 = \{\boldsymbol{A}^2, \frac{28}{33}\boldsymbol{A}^2 + \frac{5}{33}\boldsymbol{A}^1\} \\ \mathcal{N}_3 = \{\frac{77}{100}\boldsymbol{A}^1 + \frac{23}{100}\boldsymbol{A}^2, \frac{1}{2}\boldsymbol{A}^1 + \frac{1}{2}\boldsymbol{A}^2\}; & \mathcal{N}_4 = \{\frac{23}{100}\boldsymbol{A}^1 + \frac{77}{100}\boldsymbol{A}^2, \frac{1}{2}\boldsymbol{A}^1 + \frac{1}{2}\boldsymbol{A}^2\}. \end{array}$$

It is easy to check that these four necessary conditions are not simultaneously satisfiable by any partition with only three (convex) regions. Indeed, at least 5 are required. The η at which the necessary conditions corresponding to the extreme points of $\mathcal{C}_\eta$ provide a certificate that at least 5 regions are required for any partition to achieve a η-improvement or greater, is $\hat{\eta} \approx 27/7 - 3.2770 \approx 0.5801$. The corresponding value of λ is $\hat{\lambda} \approx 0.797$. Thus, examining values of $\eta \in [0, \hat{\eta}]$, the necessary conditions implied by the extreme points of $\mathcal{C}_\eta$ are sufficient to reveal that two-adaptability is no better than the robust solution, and in addition, they reveal the limit of 3-adaptability. Furthermore, they reveal the optimal 3-partition to be: $[0,1] = [0, \hat{\lambda}] \cup [\hat{\lambda}, 1 - \hat{\lambda}] \cup [1 - \hat{\lambda}, 1]$, for $\hat{\lambda} \approx 0.797$. Thus the value of 3-adaptability is 0.5801.

Finally, let us consider $\eta_2 = (27/7) - 2.9$. In this case, we are asking for more improvement than even the re-optimization formulation could provide (recall $\text{ReOpt}(\mathcal{P}) = 3$). In short, such improvement is not possible within our framework of a deterministic adversary. Proposition 3 tells us how the polytope $\mathcal{C}_{\eta_2}$ and the set $\mathcal{A}_{\eta_2}$ witness this impossibility. The polytope $\mathcal{C}_{\eta_2}$ has 31 vertices. It is enough to consider one of these vertices in particular: $\boldsymbol{v} = ((9/10, 1/10, 9/5), (0,0,0))$. The corresponding necessary condition is: $\mathcal{N} = \{\boldsymbol{A}^1, \boldsymbol{A}^1, \boldsymbol{A}^1\}$. Evidently, no number of partitions can ever satisfy this necessary condition. Indeed, this is precisely what Proposition 3 says: if progress η is not possible, it must be because $\bar{\mathcal{A}}_\eta \cap \mathcal{P} \neq \emptyset$.

Consider now the affine adaptability proposal of Ben-Tal et al. ([2]), and denote the optimal affine adaptability function by $\boldsymbol{x}^L(\lambda)$. The third component, $x_3^L(\lambda)$ must satisfy: $x_3^L(0), x_3^L(1) \geq 1$. Therefore by linearity, we must have $x_3^L(\lambda) \geq 1$ for all $\lambda \in [0,1]$. Furthermore, for $\lambda = 1/2$, we must also have

$$\frac{1}{2}\left(\frac{1}{2}x_1^{\text{aff}}(1/2) + \frac{1}{5}x_2^{\text{aff}}(1/2)\right) + \frac{1}{2}\left(\frac{1}{5}x_1^{\text{aff}}(1/2) + \frac{1}{2}x_2^{\text{aff}}(1/2)\right) \geq 1,$$

which implies, in particular, that $x_1^{\text{aff}}(1/2) + x_2^{\text{aff}}(1/2) \geq \frac{20}{7}$. The cost obtained by affine adaptability is

$$\begin{array}{rrl} \text{Affine}(\mathcal{P}) = & \max: & x_1^{\text{aff}}(\lambda) + x_2^{\text{aff}}(\lambda) + x_3^{\text{aff}}(\lambda) \\ & \text{s.t.}: & \lambda \in [0,1]. \end{array}$$

This is at least the value at $\lambda = 1/2$. But this is: $x_1^{\text{aff}}(1/2) + x_2^{\text{aff}}(1/2) + x_3^{\text{aff}}(1/2) \geq 20/7 + 1 = 27/7$, which is the robust value. Therefore in this case, affine adaptability is no better than the robust value.

[1]These computations were done using the software CDD by Komei Fukuda. This is an implementation of the double description method. See also http://www.cs.mcgill.ca/ fukuda/soft/cddman/node2.html for further details.

5 Right Hand Side Uncertainty and Complexity

While matrix uncertainty subsumes right hand side (RHS) uncertainty, we can say more for the special case of RHS-uncertainty. We note that demand uncertainty, capacity uncertainty, and supply uncertainty, are all modeled through RHS-robustness. We show that finding the best partition of the uncertainty set into two, can be cast as a combinatorial optimization problem. For the case where $\mathcal{P}$ has the structure of a simplex, we obtain a mixed integer LP with $\{0,1\}$-variables. We further show that optimally computing k-adaptability is NP-hard. The robust problem and LP formulation are given by:

$$\left[\begin{array}{rll} \min: & \boldsymbol{c}'\boldsymbol{x} & \\ \text{s.t.}: & \boldsymbol{A}\boldsymbol{x} \geq \boldsymbol{b}, & \forall \boldsymbol{b} \in \mathcal{P} \\ & \boldsymbol{D}\boldsymbol{x} \geq \boldsymbol{d} & \end{array}\right] = \left[\begin{array}{rl} \min: & \boldsymbol{c}'\boldsymbol{x} \\ \text{s.t.}: & \boldsymbol{A}\boldsymbol{x} \geq \boldsymbol{b}_R \\ & \boldsymbol{D}\boldsymbol{x} \geq \boldsymbol{d} \end{array}\right]$$

where $\boldsymbol{b}_R$ is the point whose i^{th} coordinate is the maximization over $\mathcal{P}$ in the i^{th} coordinate direction. The smallest hypercube containing $\mathcal{P}$ corresponds to the single point $\boldsymbol{b}_R$: $(\mathcal{P})_R = \{\boldsymbol{b} : \boldsymbol{b} \leq \boldsymbol{b}_R\}$ (lower bounds do not affect the optimization). As we did for the general robust problem, we consider pairs of hypercubes containing $\mathcal{P}$, rather than just a single hypercube. We formalize this in the following.

Definition 1 *We say that a pair of points $(\boldsymbol{b}_1, \boldsymbol{b}_2)$* **covers** *the set $\mathcal{P}$ if for any $\boldsymbol{b} \in \mathcal{P}$ we have $\boldsymbol{b} \leq \boldsymbol{b}_1$ or $\boldsymbol{b} \leq \boldsymbol{b}_2$. We denote by $\mathcal{C}(\mathcal{P})$ the set of all pairs $(\boldsymbol{b}_1, \boldsymbol{b}_2)$ that cover the set $\mathcal{P}$:*

$$\mathcal{C}(\mathcal{P}) \stackrel{\triangle}{=} \{(\boldsymbol{b}_1, \boldsymbol{b}_2) : \forall \boldsymbol{b} \in \mathcal{P} \text{ we have } \boldsymbol{b} \leq \boldsymbol{b}_1, \text{ or } \boldsymbol{b} \leq \boldsymbol{b}_2\}.$$

Next, we define the problem:

$$\begin{array}{rcrl} z_2(\boldsymbol{b}_1, \boldsymbol{b}_2) & = & \min: & \boldsymbol{c}'\boldsymbol{x}_1 \vee \boldsymbol{c}'\boldsymbol{x}_2 \\ & & \text{s.t.}: & \boldsymbol{A}\boldsymbol{x}_1 \geq \boldsymbol{b}_1 \\ & & & \boldsymbol{A}\boldsymbol{x}_2 \geq \boldsymbol{b}_2 \\ & & & \boldsymbol{D}\boldsymbol{x}_l \geq \boldsymbol{d}, \quad l = 1, 2. \end{array}$$

Thus, we can rewrite the problem of obtaining the best partition as the optimization of the convex function $z_2(\boldsymbol{b}_1, \boldsymbol{b}_2)$ over the set $\mathcal{C}(\mathcal{P})$: $\min : \{z_2(\boldsymbol{b}_1, \boldsymbol{b}_2) \,|\, (\boldsymbol{b}_1, \boldsymbol{b}_2) \in \mathcal{C}(\mathcal{P})\}$.

Proposition 5 *The set $\mathcal{C}(\mathcal{P})$ can be written as a finite union of polyhedral sets.*

Consider the set

$$\mathcal{C}(\mathcal{P}, S) \stackrel{\triangle}{=} \{(\boldsymbol{b}_1, \boldsymbol{b}_2) \in \mathcal{C}(\mathcal{P}) : (\boldsymbol{b}_1)_i = (\boldsymbol{b}_R)_i, \forall i \in S, \quad (\boldsymbol{b}_2)_j = (\boldsymbol{b}_R)_j, \forall j \in S^c\}.$$

The union of the sets $\mathcal{C}(\mathcal{P}, S)$ for all S, gives us back the subset of $\mathcal{C}(\mathcal{P})$ that always contains the optimal solution. We can characterize the sets $\mathcal{C}(\mathcal{P}, S)$ using a linear optimization approach.

Lemma 3 *For any fixed set $S \subseteq \{1, \ldots, m\}$, the set $\mathcal{C}(\mathcal{P}, S)$ is given as the set of all pairs $(\boldsymbol{b}_1, \boldsymbol{b}_2)$ with*

$$(\boldsymbol{b}_1)_i = \begin{cases} (\boldsymbol{b}_R)_i & \text{if } i \in S, \\ \mu_i, & \text{otherwise,} \end{cases} \quad \text{and} \quad (\boldsymbol{b}_2)_j = \begin{cases} (\boldsymbol{b}_R)_j & \text{if } j \in S^c, \\ \lambda_j, & \text{otherwise.} \end{cases}$$

The values $\{\mu_i\}$ must satisfy $0 \leq \mu_i \leq (\boldsymbol{b}_R)_i$. The values $\{\lambda_j\}$ must be such that:

$$\lambda_i \geq \max_{j \in S^c} \left[\begin{array}{rl} \max: & \boldsymbol{e}_i'\boldsymbol{b} \\ \text{s.t.}: & (\boldsymbol{b})_j \geq \mu_j \\ & \boldsymbol{b} \in \mathcal{P} \end{array}\right].$$

In an important special case, the lemma above says that the sets $\mathcal{C}(\mathcal{P}, S)$ are polyhedral. Therefore if we can compute the optimal set S^*, then computing the optimal pair $(\boldsymbol{b}_1, \boldsymbol{b}_2) \in \mathcal{C}(\mathcal{P}, S^*)$ amounts to solving an LP.

Proposition 6 *(1) Let $\mathcal{P}$ be a generalized simplex, that is, $\mathcal{P} = \text{conv}\{\zeta_1 \boldsymbol{e}_1, \zeta_2 \boldsymbol{e}_2, \ldots, \zeta_m \boldsymbol{e}_m\}$, where $\boldsymbol{e}_i$ is the standard i^{th} unit basis vector, and the ζ_i are positive scalars. Then, for any $S \subseteq \{1, \ldots, m\}$ the set $\mathcal{C}(\mathcal{P}, S)$ is convex, and polyhedral, given by:*

$$\mathcal{C}(\mathcal{P}, S) = \left\{ (\boldsymbol{b}_1, \boldsymbol{b}_2) \ : \ \begin{array}{ll} (\boldsymbol{b}_1)_i = (\boldsymbol{b}_R)_i = \zeta_i & \forall i \in S \\ (\boldsymbol{b}_2)_j = (\boldsymbol{b}_R)_j = \zeta_j & \forall j \in S^c \\ (\boldsymbol{b}_2)_i \geq \zeta_i - \left(\frac{\zeta_i}{\zeta_j}\right)(\boldsymbol{b}_1)_j & \forall i \in S, \ j \in S^c \end{array} \right\}.$$

(2) For $\mathcal{P}$ as above, the optimal split and corresponding contingency plans, may be computed as the solution to a $\{0,1\}$ mixed integer linear program.

Proposition 7 *Obtaining the optimal partition $\mathcal{P} = \mathcal{P}_1 \cup \mathcal{P}_2$ is in general NP-hard.*

In particular, computing 2-adaptability is NP-hard. We obtain our hardness result using a reduction from Partition, which is NP-complete ([12]).
Proof. The data for the Partition problem are the positive numbers $v_1, \ldots, v_m$. The problem is to minimize $|\sum_{i \in S} v_i - \sum_{j \in S^c} v_j|$ over subsets S. Given any such collection of numbers, consider the polytope $\mathcal{P} = \text{conv}\{\boldsymbol{e}_1 v_1, \ldots, \boldsymbol{e}_m v_m\}$. Thus, $\mathcal{P}$ is the simplex in $\mathbb{R}^m$, but with general intercepts v_i. Consider the robust optimization problem: min : $\{\sum_i x_i \,|\, I\boldsymbol{x} \geq \boldsymbol{b}, \quad \forall \boldsymbol{b} \in \mathcal{P}\}$. Suppose the optimal partition is $\mathcal{P} = \mathcal{P}_1 \cup \mathcal{P}_2$. Then we can show (see [5]) that the corresponding optimal value is

$$\frac{((v_1 + \cdots v_k) + (v_{k+1} + \cdots + v_m))^2 - (v_1 + \cdots + v_k)(v_{k+1} + \cdots + v_m)}{(v_1 + \cdots + v_m)}.$$

The first term in the numerator, and also the denominator, are invariant under choice of partition. Thus the second term in the numerator must be maximized. This is equivalent to maximizing the product $\left(\sum_{i \in S} v_i\right)\left(\sum_{j \in S^c} v_j\right)$ over $S \subseteq \{1, \ldots, m\}$, which is equivalent to the Partition problem. □

6 Efficient Algorithms

In Section 5 we establish that even 2-adaptability is NP-hard. In this section we propose a heuristic tractable algorithm. The algorithm is based on the following (see [5]):

Lemma 4 *Consider the set of partitions $\mathcal{P} = \mathcal{P}_1 \cup \mathcal{P}_2$ given by a hyperplane division of $\mathcal{P}$. If the orientation (i.e., the normal vector) of the hyperplane is given, then selecting the optimal partitioning hyperplane with this normal can be done efficiently.*

Algorithm 1: Let the uncertainty region be given as a convex hull of extreme points: $\mathcal{P} = \text{conv}\{\boldsymbol{A}^1, \ldots, \boldsymbol{A}^K\}$.
(1) For every pair (i, j), $1 \leq i \neq j \leq K$, let $\boldsymbol{v}_{ij} = \boldsymbol{A}^j - \boldsymbol{A}^i$.
(2) Consider the family of hyperplanes with normal $\boldsymbol{v}_{ij}$.
(3) Solve the quasi-convex problem, and let $\boldsymbol{H}_{ij}$ be the hyperplane that defines the optimal hyperplane partition of $\mathcal{P}$ within this family.
(4) Select the optimal (i, j) and the corresponding optimal hyperplane partition of $\mathcal{P}$.

This algorithm can be applied iteratively, as a heuristic approach to computing 2^d-adaptability. In [5] we show that Section 3.2 gives a structured approach to increasing the size of the family of partitions considered by introducing interior points of $\mathcal{P}$ that allow us to consider a larger collection of normal directions. We also give a counterpart algorithm for the case where $\mathcal{P}$ is given by inequalities.

7 Computational Examples

We illustrate the k-adaptability proposal and the above algorithm by some small examples.

Network Flow with Uncertain Capacity Constraints

We consider a network flow problem subject to capacity uncertainty. This highly simplified problem comes from a model of air-traffic control where capacity is affected by an arriving storm front (see [8] for a full model). The exact time and location where a storm front affects an air traffic corridor are subject to some uncertainty. In Figure 1 we show a simple network. All edges have capacity 20. Edges 1,4,8 and 2,5,9 represent direct routes

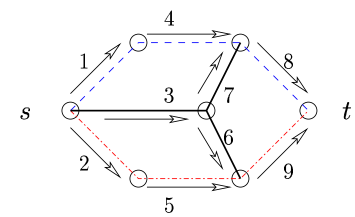

Figure 1: A small network illustrating a potential application of k-adaptability to network flows.

with cost 1. Edges 3,5,7 represent longer routes entering neighboring air space, and their cost is 6. The decision-maker must schedule 30 units of flow (airplanes) from s to t. Suppose the capacity of the central corridor marked with a solid line remains unaffected (due to geographic distance) while other capacities are affected by the storm. We consider an uncertainty region described by four extreme points, each representing some capacity degradation. The first extreme point represents loss of 24% of capacity in edge 1 and 28% in edge 8, the next 6% in edge 1 and 22% in 4, the third 36% in edge 2 and 44% in 5, and the fourth 11% in edge 5 and 12% in 9. With exact information, the flight controller can reroute flights, completely avoiding any additional cost due to the storm. Under the robust formulation, the threatening storm increases the cost by over 49%. The 2-adaptability computed by our heuristic algorithm outperforms the robust cost by over 36%. Thus, under 2-adaptability, the cost of the storm is at most 13% above the clear-weather cost, down from 49%.

Robust Parallel Scheduling

We consider a parallel scheduling problem. Let $\boldsymbol{A}$ be the rate matrix, with a_{ij} the rate of completion of product i at station j. Let c_j be the cost per hour of running station j, and let x_j be the decision variable of how long to operate station j. To minimize the operating cost subject to completing all the work, we solve min : $\{\boldsymbol{c}'\boldsymbol{x} \mid \boldsymbol{A}\boldsymbol{x} \geq 1, \boldsymbol{x} \geq 0\}$. In the robust version of the problem, the rate matrix $\boldsymbol{A}$ is only known to lie in some set $\mathcal{P}$. How much can we reduce our cost if we can formulate 2 (in general k) schedules rather than just one? Particularly in the case of binary decisions about which stations to use, preparing k contingency plans as opposed to just one, may be costly. It is therefore natural to seek to understand the value of k-adaptability, so the optimal trade-off may be selected.

We generate a large ensemble of these problems, and we report average results. The intention is to obtain some understanding for the value of adaptability, and the performance of our heuristic algorithm, in the generic case. We consider the gap between the robust problem and the re-optimization. The re-optimization is typically difficult to compute exactly ([2]). Therefore we compute upper bounds on the gap between the robust and the re-optimization values. Thus we present lower bounds on the benefit of adaptability and the performance of our algorithm. We obtain upper bounds on the gap by approximating the re-optimization value by sampling. Next, we compute the extent to which 2-adaptability, and 4-adaptability, as computed by the algorithm of Section 6, close this gap.

Table 1 reports the average gap between the robust and re-optimization values, as a fraction of the robust value: GAP = $(\text{Robust}(\mathcal{P}) - \text{ReOpt}(\mathcal{P})) / \text{Robust}(\mathcal{P})$. Then we report the average percentage of this gap covered by 2-adaptability and 4-adaptability, as computed by the heuristic algorithm.

Matrix Size	Size of $\mathcal{P}$	Avg Gap %	Avg 2-Adapt %	Avg 4-Adapt %
6×6	$K = 3$	10.10	63.22	70.70
6×6	$K = 6$	14.75	42.72	52.33
6×6	$K = 8$	18.45	39.15	47.42
10×10	$K = 3$	10.12	50.67	63.29
10×10	$K = 5$	14.22	38.58	49.36
10×10	$K = 10$	18.27	31.17	40.18
15×25	$K = 3$	8.06	39.27	54.53
15×25	$K = 5$	10.73	25.12	35.52
15×25	$K = 7$	13.15	18.21	26.84

Table 1: The matrices in these instances were generated independently. The first group of two columns identifies the size of the problem, where by Matrix size we mean the "number of products by number of stations," and by size of $\mathcal{P}$ we indicate the number of extreme points. All averages are over 50 independently generated problems.

In [5] we also consider adding integrality constraints. The heuristic algorithm proposed in Section 6 is tractable because of the quasi-convexity of the search for the optimal dividing hyperplane and by the limited set of normal directions considered. Both these factors are independent of the continuous or discrete nature of the underlying problem. Indeed, all that is required for the algorithms is a method to solve the robust problem.

References

[1] A. Atamtürk and M. Zhang. Two-stage robust network flow and design under demand uncertainty. Technical Report BCOL.04.03, IEOR, University of California–Berkeley, December 2004.

[2] A. Ben-Tal, A. Goryashko, E. Guslitzer, and A. Nemirovski. Adjustable robust solutions of uncertain linear programs. *Math. Programming*, 99:351–376, 2003.

[3] A. Ben-Tal and A. Nemirovski. Robust convex optimization. *Math. Oper. Res.*, 23:769–805, 1998.

[4] A. Ben-Tal and A. Nemirovski. Robust solutions of linear programming problems contaminated with uncertain data. *Math. Programming*, 88:411–421, 2000.

[5] D. Bertsimas and C. Caramanis. Finite adaptability in linear optimization. Technical Report available from: http://web.mit.edu/cmcaram/www/, M.I.T., September 2005.

[6] D. Bertsimas and M. Sim. Robust discrete optimization and network flows. *Mathematical Programming Series B*, 98:49–71, 2003.

[7] D. Bertsimas and M. Sim. The price of robustness. *Operations Research*, 52(1):35–53, 2004.

[8] D. Bertsimas and S. Stock. The traffic flow management rerouting problem in air traffic control: A dynamic network flow approach. *Transportation Science*, 34(3):239–255, 2000.

[9] E. Erdoğan, D. Goldfarb, and G. Iyengar. Robust portfolio management. Technical Report CORC TR-2004-11, IEOR, Columbia University, November 2004.

[10] C. Floudas and V. Visweswaran. *Quadratic Optimization*, pages 217–270. Handbook of Global Optimizaion. Kluwer Academic Publishers, 1994.

[11] L. El Ghaoui, F. Oustry, and H. Lebret. Robust solutions to uncertain semidefinite programs. *Siam J. Optimization*, 9(1), 1998.

[12] B. Korte and J. Vygen. *Combinatorial Optimization.* Springer-Verlag, 2002.

[13] H. Sherali and A. Alameddine. A new reformulation-linearization technique for bilinear programming problems. *Journal of Global Optimization*, 2:379–410, 1992.

[14] A. Takeda, S. Taguchi, and R. Tütüncü. Adjustable robust optimization models for nonlinear multi-period optimization. *Submitted*, 2004.