

Pareto Efficiency in Robust Optimization

Dan A. Iancu

Graduate School of Business, Stanford University, Stanford, California 94305, daniancu@stanford.edu

Nikolaos Trichakis

Harvard Business School, Harvard University, Boston, Massachusetts 02163, ntrichakis@hbs.edu

This paper formalizes and adapts the well-known concept of Pareto efficiency in the context of the popular robust optimization (RO) methodology for linear optimization problems. We argue that the classical RO paradigm need not produce solutions that possess the associated property of Pareto optimality, and we illustrate via examples how this could lead to inefficiencies and suboptimal performance in practice. We provide a basic theoretical characterization of Pareto robustly optimal (PRO) solutions and extend the RO framework by proposing practical methods that verify Pareto optimality and generate solutions that are PRO. Critically important, our methodology involves solving optimization problems that are of the same complexity as the underlying robust problems; hence, the potential improvements from our framework come at essentially limited extra computational cost. We perform numerical experiments drawn from three different application areas (portfolio optimization, inventory management, and project management), which demonstrate that PRO solutions have a significant potential upside compared with solutions obtained using classical RO methods.

Key words: robust optimization; Pareto optimality; linear programming

History: Received August 1, 2012; accepted March 5, 2013, by Gérard P. Cachon, optimization. Published online in *Articles in Advance* September 16, 2013.

1. Introduction

Robust optimization (RO) is a relatively young methodology, developed mainly in the course of the last 15 years to analyze and optimize the performance of complex systems; refer to the survey papers by Ben-Tal and Nemirovski (2007) and Bertsimas et al. (2011a) and the book by Ben-Tal et al. (2009) for a thorough overview of the framework. As a result of its versatility and tractability, recent years have seen an explosion of applications of RO in management science, ranging from inventory management to dynamic pricing and revenue management, portfolio optimization, and healthcare applications (for a comprehensive list, see Gabrel et al. 2012).¹ The goal of this paper is not to motivate the use of robust (and, more generally, distribution-free) techniques, but rather to introduce a methodological enhancement to this popular modeling approach.

RO deals with decision-making problems where some of the parameters are a priori unknown and/or subject to uncertainty. The standard approach is to assume that all such parameters belong to particular uncertainty sets and to take decisions so as to optimize the worst-case performance among all

possible uncertainty realizations. As such, RO is indeed indifferent toward non-worst-case scenarios and the performance of decisions made under them.

In this paper, we discuss and demonstrate via numerous examples how the inherent focus of RO on optimizing performance only under worst-case outcomes might leave decisions “unoptimized” in case a non-worst-case scenario materialized. Clearly, this is undesirable, particularly in circumstances where these decisions are implemented in practice, because they lead to suboptimal performance. To the best of our knowledge, this paper is the first to reveal this extra dimension for optimization in the classical RO framework and the first to propose a scalable, tractable method for exploiting it in a way that *strictly* enhances the framework: indeed, our approach comes with *no downside* in either performance or computational complexity.

To formalize our findings and the conclusion that RO might lead to decisions that are unoptimized for non-worst-case outcomes, we introduce the concept of Pareto efficiency in RO. The concept mimics the corresponding one in economics, engineering, and multiobjective optimization: a decision that is Pareto robustly optimal is guaranteed to deliver optimized performance across all possible scenario realizations, in the same way, for instance, that a Pareto-optimal solution in multiobjective optimization delivers optimized performance across all different objectives. We demonstrate via examples that

¹ This survey reports 100 papers authored in the areas of operations research and management science between 2007 and (June) 2012 containing the words “robust optimization” in their title, as well as a further 762 articles containing either “robust” or “robustness” in their title.

decisions made using RO need not have this property. To alleviate this, we propose methods for verifying whether a decision is Pareto optimal or not and methods for obtaining robustly optimal decisions that are provably Pareto optimal as well. Put differently, in this paper, we introduce an *essential* property that RO decisions that are to be implemented in practice need to possess (in a way that parallels the importance of Pareto efficiency in economics) and develop theory and computational tools pertaining to it. Our methodology enables a decision maker to compute robustly optimal solutions that are compatible with the classical RO framework, incur no extra computational cost, and can perform strictly better in practice.

Specifically, we make the following contributions:

a. We formalize and adapt the well-accepted concept of Pareto efficiency to the classical RO framework. We demonstrate that the framework need not produce solutions that possess the associated property of Pareto optimality, and we illustrate via examples how this could lead to inefficiencies and suboptimal performance in practice.

b. We provide a basic theoretical characterization of Pareto robustly optimal solutions.

c. We extend the RO framework by proposing practical methods that verify Pareto optimality and generate solutions that are also (provably) Pareto. Critically important, all of our proposed methodology involves optimization problems that are of *the same computational complexity* as the underlying robust problems; hence, the potential improvements from our framework come at essentially limited extra computational cost.

d. We perform numerical experiments drawn from three different application areas studied in the management science literature: portfolio optimization, inventory management, and project management. The studies demonstrate that Pareto robustly optimal solutions obtained using our methodology have a significant upside compared with solutions obtained via classical methods.

We conclude by noting that our treatment in this paper is restricted to the case of uncertain linear optimization problems. We make this choice primarily because of the overwhelming preponderance of linear models in practice, as well as for reasons of simplicity and ease of exposition. However, we see this as a first step in treating more general optimization problems appearing in the classical RO framework (Ben-Tal et al. 2009).

1.1. Literature Review

Originally introduced in operations research and management science by Soyster (1973), the RO methodology has been revitalized in the late 1990s and

early 2000s through the seminal work of Ben-Tal and Nemirovski (2002), Bertsimas and Sim (2003), and El-Ghaoui et al. (1998) (see the review papers by Ben-Tal and Nemirovski 2007 and Bertsimas et al. 2011a and the book by Ben-Tal et al. 2009 for a thorough overview of the methodology and contributions).

With a strong emphasis on computational tractability and the ability to accommodate a diverse range of relevant optimization models, the RO methodology has been adopted in many applications of interest in management science. Such examples include inventory management (Ben-Tal et al. 2005, Bertsimas and Thiele 2006, Bienstock and Özbay 2008, See and Sim 2010, Bertsimas et al. 2011b), dynamic pricing and revenue management (Adida and Perakis 2006, Lim and Shanthikumar 2007, Perakis and Roels 2010), assortment planning (Rusmevichientong and Topaloglu 2012), portfolio optimization (Goldfarb and Iyengar 2003, Fabozzi et al. 2007, Bertsimas and Pachamanova 2008), project management (Cohen et al. 2007, Adida and Joshi 2009, Goh et al. 2010, Wiesemann et al. 2012), healthcare (Chan et al. 2006), and auction design (Bandi and Bertsimas 2012). The list is by no means exhaustive; the interested reader can refer to the recent review papers by Bertsimas et al. (2011a) and Gabrel et al. (2012) and the book by Ben-Tal et al. (2009) for comprehensive references.

Despite its empirical success, however, the robust framework has been known to suffer from several potential shortcomings. One criticism is that by focusing exclusively on the worst-case outcomes, it may result in conservative decisions, with limited potential upside. This has led to several alternative proposals of robustness measures, such as absolute or relative regret (Savage 1972), “soft-robustness” (Ben-Tal et al. 2010), “light-robustness” (Fischetti and Monaci 2009), *bw*-robustness (Roy 2010, Gabrel et al. 2013), α -robustness (Kalai et al. 2012), among others. Depending on the exact assumptions and setup, such approaches typically result in the same (or slightly decreased) modeling flexibility and the same (or slightly increased) computational complexity as the standard RO framework. Critically, however, all such approaches trade off some of the robustness (i.e., performance in the worst case) in exchange for potential upside. By contrast, our approach, which parallels the notion of Pareto efficiency, guarantees the same worst-case outcome while at the same time allowing potentially improved upside, with no increase in computational complexity.

A second and more subtle criticism is that by solely focusing on worst-case outcomes, the minimax/maximin criterion may result in multiple optimal solutions, and hence it may generate Pareto inefficiencies in the decision process. This idea has emerged in several research streams, from fairness in resource allocation (Young 1995, Bertsimas et al. 2012) to

multiobjective optimization (Ogryczak 1997, Suh and Lee 2001), but it has been, to the best of our knowledge, absent from the mainstream RO literature.² The unifying characteristic in the above settings is that they are concerned with a *finite* set of alternative realizations/scenarios (e.g., the multiple parties in a problem of equitable resource allocation or the objectives in a multiobjective problem). This case, however, is typically of limited interest in RO, where the standard setup involves a continuous uncertainty set, i.e., an infinite number of realizations (Ben-Tal et al. 2009). In this sense, the proposals for correcting Pareto inefficiencies of the maximin rule in resource allocation and multiobjective optimization (e.g., the lexicographic max–min fairness scheme; see Young 1995, Ogryczak 1997) cannot be readily extended to the classical RO framework. By contrast, our definition of Pareto robustly optimal solutions is applicable to cases with both finite and infinite uncertainty sets, and hence fully compatible with the RO methodology. Furthermore, our approach still allows potential multiplicity in the Pareto robustly optimal solutions, hence not eliminating the benefits of multiplicity in RO (Bertsimas et al. 2013, Iancu et al. 2013).

Finally, we note that the concepts of “Pareto optimality/efficiency” and “robust optimization” have appeared together before, typically in the area of multiobjective optimization, but in a very different spirit than that addressed in the present paper. In particular, Gorissen and den Hertog (2011) discuss the use of RO and, more broadly, convex optimization to approximate the Pareto frontier in multiobjective optimization problems. Several papers have also attempted introducing robust formulations of multiobjective problems (e.g., Suh and Lee 2001, Deb and Gupta 2005, Cristobal et al. 2006, Luo and Zheng 2008, Ono and Nakayama 2009, Chen et al. 2012). The notions of robustness used are typically different than those in the classical RO framework, resulting in models that require solving very difficult optimization problems; this usually leads to the use of various heuristics, such as evolutionary algorithms, polynomial chaos, or particle swarm optimization.

1.2. Notation

We use $\mathbf{1}$ to denote the vector with all components equal to 1. We use e_i to denote the unit vector with 1 in the i th component. The inequality sign for vectors is used for component-wise inequality.

We use several basic notions of convex analysis (Rockafellar 1970) that we denote as follows. For a set $S \subset \mathbb{R}^n$, we use $\text{ext}(S)$ to denote the set of its extreme

points, we use $\text{conv}(S)$ to denote its convex hull, we use $\text{ri}(S)$ to denote its relative interior, and we use S^* to denote its dual cone.

2. Pareto Robustly Optimal Solutions

In this section, we introduce and formally define the notion of Pareto efficiency in RO. For illustration purposes and to ease exposition, we consider a specific form of RO problems, for which we present our definitions and the results in §3. In §4, we discuss how our findings extend to more general forms of RO problems.

The type of RO problem we consider is a linear optimization problem where only the objective is subject to uncertainty. Specifically, we consider the problem of selecting x from a polyhedral feasible set $X \subset \mathbb{R}^n$ so as to maximize $p^T x$. The objective vector p is a priori unknown and belongs to a polyhedral uncertainty set $U \subset \mathbb{R}^n$. We assume that both sets X and U are nonempty and bounded and that their inequality representations are

$$X = \{x \in \mathbb{R}^n: Ax \leq b\} \quad \text{and} \quad (1a)$$

$$U = \{p \in \mathbb{R}^n: Dp \geq d\}, \quad (1b)$$

where $A \in \mathbb{R}^{m_x \times n}$, $b \in \mathbb{R}^{m_x}$, $D \in \mathbb{R}^{m_u \times n}$, and $d \in \mathbb{R}^{m_u}$ are given. We note that our findings readily extend to the case where we only have access to a representation of U through its extreme points.

Although linear models with polyhedral uncertainty sets are a strict subset of the RO methodology, we choose them as the focus of our treatment due to their widespread use in practice, and because they are a natural first step before examining more general convex optimization models.

According to the classical RO paradigm (Ben-Tal et al. 2009), one selects x in the above setting by solving the following optimization problem:³

$$\text{maximize}_{x \in X} \min_{p \in U} p^T x. \quad (2)$$

To solve this problem, note that $\min\{p^T x: p \in U\} = \max\{y^T d: D^T y = x, y \geq 0\}$, by strong linear programming duality. Therefore, letting z^{RO} denote the optimal value and letting X^{RO} denote the set of optimal solutions of problem (2), it is easy to check that

$$X^{\text{RO}} = \{x \in X: \exists y \in \mathbb{R}_+^{m_u} \text{ such that } D^T y = x, y^T d \geq z^{\text{RO}}\}. \quad (3)$$

³Note that this model can also correspond to a distributionally robust setting, where U is a subset of the probability simplex, denoting ambiguity about the true probability measure, and the goal is to maximize the worst-case expected outcome (where nature is choosing the distribution adversarially). In this sense, our model can speak to the vast literature on risk measures and choice under ambiguity (see, e.g., Ben-Tal et al. 2010 or Bertsimas et al. 2011a for more references).

²The multiplicity of optimal solutions has been typically seen as a positive feature in RO, because it allowed deriving simple dynamic policies in multiperiod problems (Bertsimas et al. 2010, Iancu et al. 2013) or optimizing secondary objectives (Bertsimas et al. 2013).

An optimal solution $x \in X^{\text{RO}}$ of problem (2) is typically referred to as a *robustly optimal solution* in this setting. It corresponds to a solution that maximizes the worst-case objective value $p^T x$ under all possible realizations of the uncertainty $p \in U$. In other words, a robustly optimal solution is selected under the sole requirement of protecting us against worst-case scenarios. We are guaranteed that no other solution does better along that requirement. If such a solution is to be used in practice, however, this raises the following questions.

- How would x perform (in terms of the objective value) in case the uncertainty scenario that actually materialized did not correspond to a worst-case one?
- Are there any guarantees that no other solution exists that, apart from protecting us from worst-case scenarios, also performs better overall, under all possible uncertainty realizations?

To better understand the meaning of these questions, we consider a stylized network flow example arising in a communication network. The example illustrates that, in general, a robustly optimal solution $x \in X^{\text{RO}}$ might perform poorly in case a worst-case scenario was not realized: in fact, we show that it is sometimes possible to find another robustly optimal solution that outperforms x under all possible uncertainty realizations.

EXAMPLE 1. Consider a communication network that consists of multiple links. Any given link is used to transmit information between two points in the network at a rate that is determined by the network

manager. The links share different capacitated communication channels for transmission; capacity limitations then affect the information transmission rates of the individual links.

Specifically, consider the following network structure: there are two channels, denoted by A and B. The channels are of unit capacity and are utilized by $N + 3$ links, denoted by l_0, l_1, \dots, l_{N+2} , with $N \geq 3$. Links l_0 and l_1 utilize only channel A. Link l_2 utilizes both channels A and B. Links l_3, \dots, l_{N+2} utilize only channel B. Let a_i denote the transmission rate of link l_i over channel A, $i = 0, 1, 2$. Accordingly, b_i denotes the transmission rate of link l_i over channel B, $i = 2, 3, \dots, N + 2$. Vectors $a \in \mathbb{R}^3$ and $b \in \mathbb{R}^{N+1}$ contain the associated values. The structure is depicted in Figure 1.

Links l_1, \dots, l_{N+2} are dedicated for emergency purposes, whereas link l_0 is used for general purposes. Let x_i be the transmission rate of the emergency link l_i , $i = 1, 2, \dots, N + 2$. We have

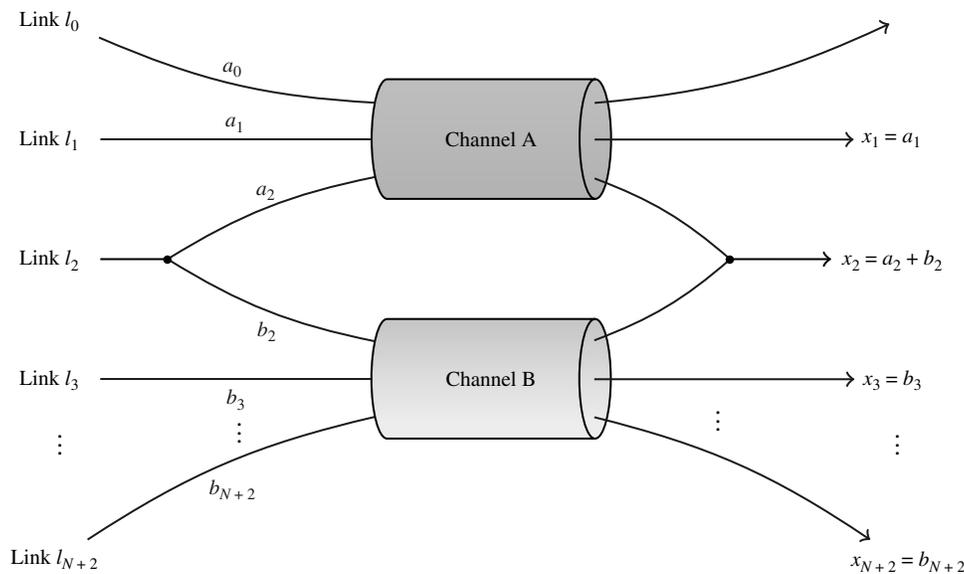
$$x_1 = a_1, \quad x_2 = a_2 + b_2, \quad \text{and}$$

$$x_i = b_i, \quad i = 3, 4, \dots, N + 2.$$

In case an emergency transmission needs to be established, this is achieved by utilizing one of the dedicated links or a combination thereof. More specifically, let f_i be the fraction of the emergency transmission routed via link l_i , $i = 1, 2, \dots, N + 2$. The net emergency transmission rate is then

$$f^T x = \sum_{i=1}^{N+2} f_i x_i.$$

Figure 1 An Illustration of the Network Structure in Example 1



Fractions f depend on the emergency situation and are uncertain. In particular, f is assumed to belong to a probability simplex uncertainty set U ; i.e.,

$$U = \{f \in \mathbb{R}_+^{N+2}: \mathbf{1}^T f = 1\}.$$

The problem for the network manager is then to select rates x , a , and b (in case of an emergency) so as to maximize the net emergency transmission rate. If the manager uses robust optimization, the rates are selected by solving

$$\begin{aligned} & \text{maximize} && \min_{f \in U} f^T x \\ & \text{subject to} && x_1 = a_1, \\ & && x_2 = a_2 + b_2, \\ & && x_i = b_i, \quad i = 3, \dots, N+2, \\ & && a_0 + a_1 + a_2 = 1, \\ & && b_2 + b_3 + \dots + b_{N+2} = 1, \\ & && a, b \geq 0, \end{aligned} \quad (4)$$

with variables x , a , and b . Let X be the feasible set. It is easy to check that the optimal value is $z^{\text{RO}} = 1/N$, and the optimal set is

$$X^{\text{RO}} = \left\{ (x, a, b) \in X: x \geq \frac{1}{N} \mathbf{1} \right\}.$$

To solve problem (4), we used the standard RO methodology introduced in §2 and solved the resulting linear program using interior point methods.⁴ For a problem with $N = 10$, we obtained the following robustly optimal solution:

$$\begin{aligned} a^{\text{IP}} &= \left[\frac{1}{3} \quad \frac{1}{3} \quad \frac{1}{3} \right]^T, & b^{\text{IP}} &= \left[0 \quad \frac{1}{10} \quad \dots \quad \frac{1}{10} \right]^T, \\ x^{\text{IP}} &= \left[\frac{1}{3} \quad \frac{1}{3} \quad \frac{1}{10} \quad \dots \quad \frac{1}{10} \right]^T. \end{aligned}$$

Note that for any realization of f drawn from U , we have

$$f^T x^{\text{IP}} \geq \frac{1}{N} = \frac{1}{10}.$$

An associated worst-case realization of f for which the obtained performance is equal to $1/10$ is, for instance, e_3 . Consider now a non-worst-case scenario for x^{IP} , e.g., $f = e_1$. Then, the realized objective value is $e_1^T x^{\text{IP}} = 1/3$.

We now compare the robustly optimal solution $(x^{\text{IP}}, a^{\text{IP}}, b^{\text{IP}})$ that we obtained with the following solution:

$$\begin{aligned} x^* &= \left[\frac{1}{2} \quad \frac{1}{2} \quad \frac{1}{10} \quad \dots \quad \frac{1}{10} \right]^T, \\ a^* &= \left[0 \quad \frac{1}{2} \quad \frac{1}{2} \right]^T, & b^* &= \left[0 \quad \frac{1}{10} \quad \dots \quad \frac{1}{10} \right]^T. \end{aligned}$$

It is easy to check that this solution is also robustly optimal; i.e., $(x^*, a^*, b^*) \in X^{\text{RO}}$. Hence, it has the same

⁴ More specifically, the SeDuMi software package (Sturm et al. 2006) was used.

qualities as the obtained solution $(x^{\text{IP}}, a^{\text{IP}}, b^{\text{IP}})$ in protecting us from worst-case realizations of f . However, if $f = e_1$ is realized, this solution yields strictly better performance:

$$e_1^T x^* = \frac{1}{2} > e_1^T x^{\text{IP}} = \frac{1}{3}.$$

In fact, it is easy to see that x^* performs strictly better or at least as well as x^{IP} for *any* realization of f drawn from U . In particular, we have

$$f^T x^* \geq f^T x^{\text{IP}}, \quad \forall f \in U, \quad \text{and} \quad (5a)$$

$$f^T x^* > f^T x^{\text{IP}}, \quad \forall f \in U \cap \{f: f_1 + f_2 > 0\}. \quad (5b)$$

The above discussion demonstrates that by focusing only on worst-case scenarios, the classical RO paradigm allows solutions that are *dominated*, i.e., that can be improved without affecting worst-case performance. Clearly, this inefficiency is not important in case RO is used only for purposes of determining the optimal worst-case costs or profits. However, it might negatively affect managers who utilize RO to actually compute solutions they implement in practice.

To rectify the aforementioned weakness of the classical RO framework, we need to additionally require from a robustly optimal solution to perform as well as possible across all uncertainty scenarios that can be realized. In particular, for any robustly optimal solution we select, we need to guarantee that there does not exist another solution that performs at least as well across all uncertainty realizations and strictly better at some realizations. If such a solution existed, it would be strictly preferred for all practical considerations.

We call solutions with the above property *Pareto robustly optimal* (PRO) solutions. Formally, we have the following definition.

DEFINITION 1. A solution x is a Pareto robustly optimal solution for problem (2) if it is robustly optimal, i.e., $x \in X^{\text{RO}}$, and there is no other $\bar{x} \in X$ such that

$$\begin{aligned} p^T \bar{x} &\geq p^T x, \quad \forall p \in U, \quad \text{and} \\ \bar{p}^T \bar{x} &> \bar{p}^T x, \quad \text{for some } \bar{p} \in U. \end{aligned}$$

In the definition above, we say that \bar{x} *Pareto dominates* x . As discussed in the introduction, the terminology we use borrows from economics and multiobjective optimization: RO can be viewed as a multiobjective optimization problem with an infinite number of objectives, one for each uncertainty scenario.

Returning to Example 1, by (5), we have that the solution (x^*, a^*, b^*) Pareto dominates the solution x^{IP} . Moreover, one can show that the solution (x^*, a^*, b^*) is

a PRO solution. In fact, if we denote the set of all PRO solutions with X^{PRO} , then for problem (4) we have

$$X^{\text{PRO}} = \{(x, a, b) \in X: x \geq \frac{1}{N}\mathbf{1}, x_1 + x_2 = 1\}.$$

Below is another toy example that illustrates the notion of PRO solutions.

EXAMPLE 2. Consider the following problem, which is of the same form as problem (2):

$$\begin{aligned} &\text{maximize} && \min_{p \in U} \{p_1x_1 + p_2x_2 + p_3x_3\} \\ &\text{subject to} && x_1 - x_2 = 0, \\ &&& x_1 + x_3 = 0, \\ &&& x_1 \geq 0, \\ &&& x_1 \leq 1, \end{aligned} \tag{6}$$

where $U = \{p \in \mathbb{R}^3: \mathbf{1} \leq p \leq \mathbf{2}\}$ is a hypercube uncertainty set.

For the above problem, it is easy to check that for any $x \in X$, $[1 \ 1 \ 2]^T$ is a worst-case uncertainty realization for which $[1 \ 1 \ 2]x = 0$. Hence, we have that $X = X^{\text{RO}}$.

To solve problem (6), we used the standard RO methodology discussed above and solved the resulting linear program using the simplex method.⁵ We obtained the robustly optimal solution $x^{\text{simplex}} = [0 \ 0 \ 0]^T$. Consider the solution $x^* = [1 \ 1 \ -1]^T$. For any uncertainty realization different from the worst-case realization we identified above, i.e., $p \neq [1 \ 1 \ 2]^T$, we have that $p^T x^* > p^T x^{\text{simplex}}$. For the worst-case realization, both solutions yield an objective value of zero. Hence, solution x^* Pareto dominates solution x^{simplex} . In fact, x^* is the only PRO solution for problem (6), so that $X^{\text{PRO}} = \{x^*\}$.

There are many interesting questions to be addressed that are theoretically and practically relevant in view of the notion of PRO solutions we have introduced. Do PRO solutions always exist, and how can we find them efficiently? When is every robustly optimal solution also a PRO solution? Can we characterize X^{PRO} ? Is it a convex set?

Apart from shedding light on the questions raised above, the rest of the paper is devoted to answering the following two central questions:

1. Given a robustly optimal solution $x \in X^{\text{RO}}$, how do we check whether x is also PRO? If it is not, how do we find a PRO solution $\tilde{x} \in X^{\text{PRO}}$ that Pareto dominates x ?

2. How do we optimize over the set of PRO solutions X^{PRO} ? From our discussion so far, it should be obvious that for practical decision making, a manager should always prefer PRO solutions. In case a manager also has a secondary objective, how can she select a PRO solution that is optimal with respect to this secondary objective?

3. Finding and Optimizing Over PRO Solutions

As suggested at the end of §2, several questions of interest can be posed concerning the set of PRO solutions. The goal of this section is to provide detailed answers to these questions. We focus our discussion here on the class of RO problems described by (2), when the feasible set X and the uncertainty set U are given by (1a) and (1b), respectively. In §4, we revisit and extend our results to several other models of interest.

3.1. Finding PRO Solutions

In current practice, a decision maker would first formulate a RO model for the particular application of interest and then seek to solve the resulting problem, hence determining a robustly optimal solution $x \in X^{\text{RO}}$. However, as suggested by the simple examples in §2, the solution x does not necessarily have to be a PRO solution.

In this context, the first question of interest is how to check whether a given $x \in X^{\text{RO}}$ is also PRO and, if not, how to find an \tilde{x} that is PRO and Pareto dominates x . The following theorem argues that both of these questions can be answered in a straightforward fashion by solving a single linear program (LP) of compact size.

THEOREM 1. Given a point $x \in X^{\text{RO}}$, consider an arbitrary point $\bar{p} \in \text{ri}(U)$ and the following optimization problem:

$$\begin{aligned} &\text{maximize}_y && \bar{p}^T y \\ &\text{subject to} && y \in U^*, \\ &&& x + y \in X. \end{aligned} \tag{7}$$

Then, either the optimal value in the problem is zero, in which case $x \in X^{\text{PRO}}$, or the optimal value is strictly positive, in which case $\tilde{x} = x + y^*$ Pareto dominates x and $\tilde{x} \in X^{\text{PRO}}$ for any optimal solution y^* .

PROOF OF THEOREM 1. Note first that $y = 0$ is always feasible in problem (7). Hence, the optimal value is nonnegative. The discussion separates into two disjoint cases.

Case 1: The optimal value is zero. Assume that $x \notin X^{\text{PRO}}$; i.e., there exists $\tilde{x} \in X$ that Pareto dominates x . In particular, (i) $p^T \tilde{x} \geq p^T x$, $\forall p \in U$, and (ii) $\exists \hat{p} \in U$ such that $\hat{p}^T \tilde{x} > \hat{p}^T x$. Without loss of generality, we can

⁵ We used the IBM ILOG CPLEX solver, with the simplex method selected.

additionally take $\hat{p} \in \text{ext}(U)$, e.g., as a vertex solution to $\max_{p \in U} (\tilde{x} - x)^T p$.

Note that $\tilde{x} - x$ is readily feasible in (7). We claim that $\bar{p}^T (\tilde{x} - x) > 0$. To prove this, recall that any $\bar{p} \in \text{ri}(U)$ has a representation as a strict convex combination of all extreme points of U (Rockafellar 1970), i.e., $\exists \lambda \in \mathbb{R}^{|\text{ext}(U)|}$ such that $\lambda > 0$, $\mathbf{1}^T \lambda = 1$, and $\bar{p} = \sum_{i \in \mathcal{J}} \lambda_i p_i$, where $\text{ext}(U) \stackrel{\text{def}}{=} \{p_i : i \in \mathcal{J}\}$. But then,

$$\bar{p}^T (\tilde{x} - x) = \underbrace{\hat{p}^T (\tilde{x} - x)}_{>0} + \sum_{i \in \mathcal{J}: p_i \neq \hat{p}} \underbrace{p_i^T (\tilde{x} - x)}_{\geq 0} > 0.$$

This immediately leads to the desired contradiction.

Case 2: The optimal value is positive. The conditions in Definition 1 can be directly verified for \tilde{x} , showing that \tilde{x} Pareto dominates x . To see that $\tilde{x} \in X^{\text{PRO}}$, consider testing this by solving problem (7), i.e., with x replaced by $\tilde{x} = x + y^*$, and the same \bar{p} . We claim that the optimal solution to this problem must be zero (which, in turn, implies that $\tilde{x} \in X^{\text{PRO}}$). Otherwise, if an optimal solution \tilde{y} existed, with $\bar{p}^T \tilde{y} > 0$, then $y^* + \tilde{y}$ would be feasible and provide a higher objective value than y^* when solving problem (7) to test whether x was PRO, contradicting the fact that y^* was an optimal solution. \square

Let us comment on the result and its relevance. We claimed that problem (7) is an LP. This follows since the dual cone U^* of a polyhedral set U always has a polyhedral representation (Rockafellar 1970). For instance, if U is the polytope given by (1b), then by strong LP duality,

$$\begin{aligned} U^* &\stackrel{\text{def}}{=} \{y \in \mathbb{R}^n : y^T p \geq 0, \forall p \in U\} \\ &= \{y \in \mathbb{R}^n : \exists \lambda \in \mathbb{R}_+^{m_U} \text{ such that} \\ &\quad D^T \lambda = y, d^T \lambda \geq 0\}. \end{aligned} \quad (8)$$

This implies that problem (7) can be solved efficiently using standard software (e.g., CPLEX, Gurobi, SeDuMi) for sizes that are typical in real-world applications.

Note also that our result is stated for an arbitrary point \bar{p} in the relative interior of U , so that problem (7) must only be solved *once*. Since finding a point in the relative interior of a polyhedron can be done efficiently by LP techniques (Schrijver 2000), this readily leads to an efficient procedure for testing whether $x \in X^{\text{PRO}}$ and (if not) for producing points that Pareto dominate x .

In a different sense, Theorem 1 also confirms that PRO solutions to any robust LP problem always exist and suggests the following procedure for finding them.

COROLLARY 1. Consider an arbitrary point $\bar{p} \in \text{ri}(U)$. Then, all the optimal solutions to the problem $\max_{x \in X^{\text{RO}}} \bar{p}^T x$ are PRO.

The proof follows analogously to that of Theorem 1 and is omitted. It is interesting to note that finding an $x \in X^{\text{PRO}}$ is not substantially more difficult than finding an $x \in X^{\text{RO}}$ —one needs to sample a point \bar{p} from $\text{ri}(U)$, which can be done by an LP, and to solve an additional LP over the set X^{RO} . That is, although a limited additional overhead is incurred, the computational complexity of the underlying RO framework is fundamentally preserved. Moreover, since the potential gains from working with PRO solutions can be substantial (as illustrated in our examples in §5), this suggests that the framework can have considerable promise in practice.

The procedure introduced in Corollary 1 also suggests a simple way to generate (potentially different) solutions in X^{PRO} , by (i) sampling different values \bar{p} from $\text{ri}(U)$ and (ii) solving the corresponding LP over X^{RO} . The following result confirms that *all* the points of X^{PRO} can, in fact, be generated by such a procedure.

PROPOSITION 1. For any $x \in X^{\text{PRO}}$, there exists $\bar{p} \in \text{ri}(U)$ such that $x \in \arg \max_{y \in X^{\text{RO}}} \bar{p}^T y$.

PROOF OF PROPOSITION 1. Let $P \in \mathbb{R}^{n \times |\text{ext}(U)|}$ denote the matrix with columns given by the extreme points of U , and without loss of generality, assume that the last $|\text{ext}(U)| - k$ columns (where $k \leq |\text{ext}(U)|$) correspond to *all* the points $p \in \text{ext}(U)$ such that $p^T x = z^{\text{RO}} \forall x \in X^{\text{RO}}$. Also, let $\bar{P} \in \mathbb{R}^{n \times k}$ denote the matrix obtained by keeping the first k columns of P .

Consider any point $x \in X^{\text{PRO}}$. To construct the desired $\bar{p} \in \text{ri}(U)$, we note that it is enough to show that $\exists \lambda \in \mathbb{R}^k$ such that $\lambda \geq \mathbf{1}$ and $x \in \arg \max_{y \in X^{\text{RO}}} y^T \bar{P} \lambda$. The reason is that any such λ can be extended into

$$\tilde{\lambda} \stackrel{\text{def}}{=} \frac{1}{\mathbf{1}^T \lambda + |\text{ext}(U)| - k} \begin{bmatrix} \lambda \\ \mathbf{1} \end{bmatrix},$$

which satisfies the relations $\tilde{\lambda} > 0$, $\mathbf{1}^T \tilde{\lambda} = 1$, and $x \in \arg \max_{y \in X^{\text{RO}}} y^T P \tilde{\lambda}$ if and only if

$$x \in \arg \max_{y \in X^{\text{RO}}} y^T P \begin{bmatrix} \lambda \\ 0 \end{bmatrix}.$$

To this end, assume (for the purposes of deriving a contradiction) that $\forall \lambda \geq \mathbf{1}$, we have $x \notin \arg \max_{y \in X^{\text{RO}}} y^T \bar{P} \lambda$. This implies that

$$\forall \lambda \geq \mathbf{1}, \exists y(\lambda) \in X^{\text{RO}} : (y(\lambda) - x)^T \bar{P} \lambda > 0 \Rightarrow$$

$$\min_{\lambda \geq \mathbf{1}} \max_{y \in X^{\text{RO}}} (y - x)^T \bar{P} \lambda > 0 \stackrel{(*)}{\Leftrightarrow}$$

$$\max_{y \in X^{\text{RO}}} \min_{\lambda \geq \mathbf{1}} (y - x)^T \bar{P} \lambda > 0 \stackrel{(**)}{\Rightarrow}$$

$$\exists y \in X^{\text{RO}} : \max\{\mathbf{1}^T \mu : \mu^T \leq (y - x)^T \bar{P}, \mu \geq 0\} > 0.$$

Step (*) follows from the minimax theorem in convex analysis, which is applicable here since the function $(y - x)^T \bar{p} \lambda$ is bilinear in y and λ , and the set X^{RO} is compact (Rockafellar 1970). Step (**) follows from strong LP duality, which holds for the given y , since the minimization over λ is feasible and bounded below. The last step implies that $x \notin X^{\text{PRO}}$, since y Pareto dominates x . \square

The previous results also allow formulating conditions under which all the points in X^{RO} are actually PRO solutions. This is relevant because a decision maker would then not have to worry about Pareto domination—he could simply utilize any solution obtained by solving a standard robust problem. The next result, whose proof also follows directly from Theorem 1, argues that this question can also be answered by solving a compact LP.

COROLLARY 2. *Consider any $\bar{p} \in \text{ri}(U)$ and the following optimization problem:*

$$\begin{aligned} & \underset{x, y}{\text{maximize}} && \bar{p}^T y \\ & \text{subject to} && x \in X^{\text{RO}}, \\ & && y \in U^*, \\ & && x + y \in X. \end{aligned} \tag{9}$$

Then, $X^{\text{PRO}} = X^{\text{RO}}$ if and only if the optimal value is zero.

PROOF OF COROLLARY 2. Let x^* , y^* be optimal solutions of problem (9). It can be checked that if the optimal value is positive, then $x^* + y^*$ is PRO and Pareto dominates $x^* \in X^{\text{RO}}$, implying that $X^{\text{PRO}} \neq X^{\text{RO}}$. The reverse direction follows since an optimal value of zero implies that any $x \in X^{\text{RO}}$ is PRO, according to Theorem 1. \square

In practice, it may be relevant to look for simpler conditions guaranteeing $X^{\text{PRO}} = X^{\text{RO}}$. By Corollary 2, one such example is when $0 \in \text{ri}(U)$ —in this case, any $x \in X^{\text{RO}}$ yields the same objective value of zero so that $X^{\text{PRO}} = X^{\text{RO}}$. It is important to note that this condition applies to the particular objective used in the robust problem (2), i.e., $p^T x$, where p is uncertain. Clearly, by suitably translating any uncertainty set U with a nonempty relative interior, one may obtain a new model with objective $r^T x + p^T x$, where r is known exactly, $p \in \tilde{U}$ is uncertain, and $0 \in \text{ri}(\tilde{U})$. For such a model, $0 \in \text{ri}(\tilde{U})$ no longer guarantees that $X^{\text{PRO}} = X^{\text{RO}}$. Although having zero in the relative interior may occur in some cases, many of the RO models considered in applications do not typically satisfy the condition (see our numerical studies in §5 and Ben-Tal et al. 2009 for details). In fact, the condition should never be satisfied if the physical nature of an uncertain parameter prevented it from switching signs, e.g., if it corresponded to an uncertain yield in a production process, a probability, or a customer demand.

3.2. Optimizing Over the Set of PRO Solutions

In the previous section, we discussed procedures for testing and generating PRO solutions for the class of RO problems described by (2). Note that in this class of problems, the manager wishes to optimize with respect to a particular objective p ; the fact that this objective vector is uncertain gives rise to the employment of the RO methodology.

In some cases in practice, however, a manager may also have a secondary objective in mind. For instance, in scheduling problems, managers often choose to minimize the number of late jobs as their primary objective. Among all solutions that optimize this objective, managers typically prefer solutions that minimize total completion time, which is then their secondary objective (see, e.g., Leung et al. 2010). In this section, we deal with the problem of optimizing such linear secondary objectives over the set of PRO solutions.⁶

More formally, for a given objective vector $r \in \mathbb{R}^n$, we are interested in solving the following problem:

$$\begin{aligned} & \text{maximize} && r^T x \\ & \text{subject to} && x \in X^{\text{PRO}}. \end{aligned}$$

Note that if the objective lies in the relative interior of U , i.e., $r \in \text{ri}(U)$, Corollary 1 provides a direct solution: the manager can simply maximize $r^T x$ over the set X^{RO} . More broadly, however, understanding the structure and properties of the set X^{PRO} becomes relevant. The following example confirms that X^{PRO} is, unfortunately, not a convex set in general.

EXAMPLE 3 (NONCONVEXITY OF X^{PRO}). Consider the following feasible set X and uncertainty set U :

$$\begin{aligned} X &= \{x \in \mathbb{R}_+^4 : x_1 \leq 1, x_2 + x_3 \leq 6, x_3 + x_4 \leq 5, x_2 + x_4 \leq 5\}, \\ U &= \text{conv}(\{e_i, i \in \{1, \dots, 4\}\}). \end{aligned}$$

It can be checked that $z^{\text{RO}} = 1$, and $X^{\text{RO}} = \{x \in X : x \geq 1\}$. Also, $x^1 = [1 \ 2 \ 4 \ 1]^T$ and $x^2 = [1 \ 4 \ 2 \ 1]^T$ are both PRO solutions (they are the optimal solutions to the problems of maximizing $[\epsilon \ \epsilon \ 1 - 3\epsilon \ \epsilon]^T$ and $[\epsilon \ 1 - 3\epsilon \ \epsilon \ \epsilon]^T$ over X^{RO} , respectively, for some small $\epsilon > 0$). However, the point $0.5x^1 + 0.5x^2 = [1 \ 3 \ 3 \ 1]^T \notin X^{\text{PRO}}$, since it is Pareto dominated by $[1 \ 3 \ 3 \ 2]^T \in X^{\text{RO}}$.

The nonconvex structure of X^{PRO} suggests that solving optimization problems over the set may be computationally challenging in general. One particular case when this is simple is when $X^{\text{RO}} = X^{\text{PRO}}$, which can be tested using the LP in Corollary 2. When $X^{\text{RO}} \neq X^{\text{PRO}}$, the following result provides some intuition about the structure of the latter set.

⁶Note that in the classical RO framework, optimizing a linear secondary objective over X^{RO} can be cast as an LP, utilizing the polyhedral description of X^{RO} in (3).

PROPOSITION 2. If $X^{RO} \neq X^{PRO}$, then $X^{PRO} \cap \text{ri}(X^{RO}) = \emptyset$.

PROOF OF PROPOSITION 2. We first argue that if $X^{RO} \neq X^{PRO}$, then there does not exist a $\bar{p} \in \text{ri}(U)$ such that $\bar{p}^T x$ is constant for all $x \in X^{RO}$. To see this, note that if such a \bar{p} existed, then the optimal objective function in problem (7), for any point $x \in X^{RO}$, would have to be zero, implying that $X^{RO} = X^{PRO}$.

Assume now that $X^{PRO} \cap \text{ri}(X^{RO}) \neq \emptyset$, and consider any x in the intersection. By Proposition 1, there must exist $\bar{p} \in \text{ri}(U)$ such that $x \in \arg \max_{x \in X^{RO}} \bar{p}^T x$. Furthermore, since $\bar{p}^T x$ is not constant over X^{RO} , there must exist a $y \in X^{RO}$ such that $\bar{p}^T y < \bar{p}^T x$. But then, since $x \in \text{ri}(X^{RO})$, there exists a small enough $\epsilon > 0$ such that $\bar{x} \stackrel{\text{def}}{=} x + \epsilon \cdot (x - y) \in X^{RO}$, for which $\bar{p}^T \bar{x} > \bar{p}^T x$ —a contradiction. \square

The previous result shows that the set X^{PRO} is either identical to X^{RO} or contained in the boundary of the latter. This result is somewhat encouraging, since it may allow a characterization of the convex hull of X^{PRO} in particular cases. We do not pursue this idea further in the present paper.

In a different sense, the results in Proposition 2 also suggest that whether a solution $x \in X^{RO}$ is actually PRO critically depends on the algorithm used for solving the nominal RO problem. In particular, if $X^{RO} \neq X^{PRO}$ and the nominal RO problem is solved using a typical interior point method (i.e., one that returns solutions in the interior of the optimal face; see Ye 1992), then, by Proposition 2, the resulting solution $x \in X^{RO}$ will *not* be PRO. However, even when the simplex method is used, there is still no reason to a priori expect that the obtained solution would be PRO (see Example 2 of §2).

We now return to the main question of interest, i.e., developing a tractable procedure for optimizing a linear objective over the set X^{PRO} , and we present two different approaches for addressing it. The first involves solving a mixed-integer linear program (MILP) and is summarized in the following proposition.

PROPOSITION 3. For any $r \in \mathbb{R}^n$ and any $\bar{p} \in \text{ri}(U)$, let $(x^*, \mu^*, \eta^*, z^*) \in \mathbb{R}^n \times \mathbb{R}^{m_x} \times \mathbb{R} \times \{0, 1\}^{m_x}$ be an optimal solution of the following MILP:

$$\text{maximize}_{x, \mu, \eta, z} r^T x \tag{10a}$$

$$\text{subject to } x \in X^{RO}, \tag{10b}$$

$$\mu \leq M(1 - z), \tag{10c}$$

$$b - Ax \leq Mz, \tag{10d}$$

$$DA^T \mu - d\eta \geq D\bar{p}, \tag{10e}$$

$$\mu \geq 0, \eta \geq 0, z \in \{0, 1\}^{m_x}, \tag{10f}$$

where M is a sufficiently large value. Then, x^* is an optimal solution of the problem

$$\text{maximize}_{x \in X^{PRO}} r^T x.$$

PROOF OF PROPOSITION 3. We need to show that $x \in X^{RO}$ is PRO if and only if there exist $\mu \in \mathbb{R}^{m_x}$, $\eta \in \mathbb{R}$, and $z \in \{0, 1\}^{m_x}$ that satisfy (10c)–(10f).

Fix a solution $x \in X^{RO}$. By Theorem 1, x is PRO if and only if, for an arbitrarily chosen $\bar{p} \in \text{ri}(U)$, the optimal value in the problem $\max_y \{\bar{p}^T y : y \in U^*, x + y \in X\}$ is zero. Equivalently, x is PRO if and only if the optimal value in the following primal-dual LP pair is zero:

$$\begin{array}{ll} \text{maximize}_{\lambda} \bar{p}^T D^T \lambda & \text{minimize}_{\mu, \eta} \mu^T (b - Ax) \\ \text{subject to} & \text{subject to} \\ d^T \lambda \geq 0 & DA^T \mu - d\eta \geq D\bar{p} \\ AD^T \lambda \leq b - Ax & \mu \geq 0 \\ \lambda \geq 0 & \eta \geq 0. \end{array}$$

To arrive at the primal (maximization), we used the description in (1a) for the set X , the expression in (8) for the dual cone U^* , and we eliminated the auxiliary variables y . Note that the primal is trivially feasible, with the choice $\lambda = 0$ resulting in an objective value of zero. As such, whenever its optimal value is less than or equal to zero, strong LP duality must hold.

We then have that the optimal value is zero if and only if there exist $\mu \in \mathbb{R}^{m_x}$ and $\eta \in \mathbb{R}$ satisfying the constraints of the dual, such that $\mu^T (b - Ax) = 0$. This is a bilinear constraint in variables x and μ . In fact, since $\mu \geq 0$ and $b - Ax \geq 0, \forall x \in X^{RO}$, it is equivalent to the linear complementarity constraints $\mu_i \cdot (b - Ax)_i = 0, i = 1, \dots, m_x$ (Luo et al. 1996). The latter constraint can be modeled using variables $z \in \{0, 1\}^{m_x}$ and constraints (10c) and (10d), where M is sufficiently large.⁷ \square

The above approach should be very relevant in practice, since large-scale MILPs can be solved in a matter of seconds using commercially available software such as CPLEX or Gurobi. In case this is still an onerous task, one can also resort to the following simple heuristic for optimizing linear functions over X^{PRO} : solve the problem $\max_{x \in X^{RO}} \bar{p}^T x$ for (several) randomly sampled points $\bar{p} \in \text{ri}(U)$, collect all the optimal solutions in a set \hat{X}^{PRO} , and then solve the problem $\max_{x \in \hat{X}^{PRO}} r^T x$.

Because, by Corollary 1 and Proposition 1 all the PRO solutions can be generated by sampling points

⁷Note that we need $M \geq \max_{i=1, \dots, m_x} \max(\mu_i, \max_{x \in X} e_i^T (b - Ax))$. The latter term can always be bounded since X is compact. The former term may also be bounded, depending on the dual feasible set. In practice, one may simply choose an increasing sequence of M , stopping when the constraints for $\mu \leq M$ are no longer binding.

\bar{p} from the relative interior of U , the set \hat{X}^{PRO} will be a subset of the true X^{PRO} , so that this algorithm will always produce a lower bound to the true optimal value. However, coupled with a suitable upper bound (e.g., obtained by maximizing $r^T x$ over X^{RO}), the heuristic may prove satisfactory in cases where the exact MILP in Proposition 3 is difficult to solve.

4. Generalizations

In this section, we discuss several directions for extending our earlier framework and results.

4.1. Uncertainty in the Constraints

Consider a linear optimization problem where the coefficients of the constraints matrix are uncertain. Specifically, consider the following RO formulation:

$$\begin{aligned} & \text{minimize } c^T x \\ & \text{subject to } Ax \geq b, \quad \forall A \in U_A, \end{aligned} \quad (11)$$

where $c \in \mathbb{R}^n$ and $b \in \mathbb{R}^m$ are given, and $U_A \subset \mathbb{R}^{m \times n}$ is a bounded polyhedral uncertainty set of the form typically considered in the RO literature⁸ (Ben-Tal et al. 2009). Problems with uncertainty in the vectors c and b can be readily reformulated so as to have the same form as problem (11). For instance, in case c is uncertain, one can consider an equivalent epigraph formulation where the objective is replaced by an additional variable t and the extra constraint $t - c^T x \geq 0$ is added.

Let X^{RO} be the optimal set and let z^{RO} be the optimal value of problem (11). The vector of slacks for constraints in (11) for a solution x and uncertainty realization A , denoted by $s(x, A)$, is equal to

$$s(x, A) \stackrel{\text{def}}{=} Ax - b, \quad \forall x \in \mathbb{R}^n, A \in U_A.$$

In view of the above, one can express X^{RO} as follows:

$$X^{\text{RO}} = \{x \in \mathbb{R}^n: c^T x \leq z^{\text{RO}}, s(x, A) \geq 0, \forall A \in U_A\}. \quad (12)$$

A compact polyhedral description of X^{RO} is readily available whenever such a description is also available for U_A , which we have assumed to be the case here.

According to classical RO, a *robustly optimal solution* $x \in X^{\text{RO}}$ in this setting protects us from worst-case realizations of $A \in U_A$ by ensuring that the slacks are nonnegative. Moreover, it does so at a minimum

⁸ The typical modeling assumption is that the matrix A depends affinely on a collection of primitive uncertain quantities, i.e., $A(\xi) = \sum_{i=1}^m A_i \xi_i$, where $A_i \in \mathbb{R}^{m \times n}$ are known and $\xi \in \mathbb{R}^m$ denotes the primitive uncertain parameters, assumed to lie in an uncertainty set $\Xi \subset \mathbb{R}^m$. As such, whenever Ξ has a compact polyhedral representation, the corresponding set $U_A = \{A(\xi): \xi \in \Xi\}$ does as well (see Ben-Tal et al. 2009 for details).

cost $c^T x = z^{\text{RO}}$. That is, by selecting a robustly optimal solution, we are guaranteed that no other solution exists that yields nonnegative slacks under all uncertainty scenarios and at a lower cost than z^{RO} . However, in a similar spirit to our discussion in §2, such a selection criterion does not “optimize slacks” under all possible uncertain realizations—in fact, it fails to guarantee that no other solution exists yielding “larger” slacks and at the same cost as z^{RO} .

The above observation gives rise to the notion of PRO solutions for formulation (11). To evaluate the slacks provided by different solutions and uncertainty realizations, we introduce a *slack value vector* $v \in \mathbb{R}^m$ that quantifies the relative value of slack in each constraint.⁹ In particular, the slacks provided by solution x , under scenario A , are valued at

$$v^T s(x, A) = v^T (Ax - b), \quad \forall x \in \mathbb{R}^n, A \in U_A.$$

We have the following definition.

DEFINITION 2. A solution x is called a PRO solution for problem (11) if it is robustly optimal, i.e., $x \in X^{\text{RO}}$, and there is no other $\bar{x} \in X^{\text{RO}}$ such that

$$\begin{aligned} v^T s(\bar{x}, A) &\geq v^T s(x, A), \quad \forall A \in U_A, \quad \text{and} \\ v^T s(\bar{x}, \bar{A}) &> v^T s(x, \bar{A}), \quad \text{for some } \bar{A} \in U_A. \end{aligned}$$

Similar to our discussion in §2, in this setting PRO solutions guarantee that slacks are optimized for all uncertainty realizations $A \in U_A$. Compared with robustly optimal solutions, PRO solutions have the same qualities, i.e., they ensure feasibility at the lowest cost possible, but, in addition, they also provide potentially higher slacks (as valued through vector v). This is particularly important for the following reasons:

a. In the majority of problems of the form (11) that are derived from practical applications, one of the constraints typically involves the *actual* realized cost, e.g., because of an epigraph reformulation.¹⁰ As such, slack in that particular constraint immediately translates into lower actual realized cost. We refer the reader to our numerical studies in §§5.2–5.3 for examples.

⁹ We choose to compare two slack vectors by scalarization through the value vector $v \in \mathbb{R}^m$, a method that is common in multiobjective optimization (Boyd and Vandenberghe 2004). However, other multidimensional orderings may be more suitable in particular settings (for instance, the Lorenz order). Extending our framework to such cases is an interesting direction that we do not pursue in the present paper.

¹⁰ Note that an epigraph reformulation is necessary not only in cases where the actual cost vector is uncertain (as discussed above), but also when the cost of the problem is piecewise affine; see, for instance, Bertsimas et al. (2010).

b. Solutions that provide zero or low values of slack can more readily generate infeasibility as a result of potential misspecifications of other model parameters. For the value of slack, we refer the reader, for instance, to Joshi and Boyd (2009).

Hence, similar to the notion studied in §2, PRO solutions in this setting have no downside and have considerable potential upside.

From a methodological standpoint, the conditions of Pareto domination in Definition 2 are equivalent to

$$\begin{aligned} (A^T v)^T \bar{x} &\geq (A^T v)^T x, \quad \forall A \in U_A, \quad \text{and} \\ (\bar{A}^T v)^T \bar{x} &> (\bar{A}^T v)^T x, \quad \text{for some } \bar{A} \in U_A. \end{aligned}$$

If we let $p \stackrel{\text{def}}{=} A^T v$ and

$$U \stackrel{\text{def}}{=} \{A^T v: A \in U_A\}, \tag{13}$$

we arrive at the same formulation of Pareto efficiency as in Definition 1. Using this observation, one can extend all our findings in §3.1. For instance, the extended Corollary 1 would state the following.

COROLLARY 3. *Consider an arbitrary point $\bar{p} \in \text{ri}(U)$, where U is given by (13). Then, all the optimal solutions to maximize $\sum_{x \in X^{\text{RO}}} \bar{p}^T x$ are PRO for problem (11), where X^{RO} is given by (12).*

As such, all the remarks in §3.1 pertaining to our results and their practical relevance hold in the present setting as well.

4.2. Mixed-Integer Linear Optimization Models

A second direction for extending our results is the case when some of the decisions x are constrained to be integers. This is of great relevance in practical applications, since many models of real-world processes involve integrality constraints (see, e.g., Bertsimas and Sim 2003).

To this end, the present section is concerned with the following problem:

$$\begin{aligned} &\text{maximize} \quad \min_{x \in X} \min_{p \in U} p^T x \\ X &= \{x \in \mathbb{R}^n: Ax \leq b, x_i \in \mathbb{Z}, \forall i \in \mathcal{I}\}, \\ U &= \{p \in \mathbb{R}^n: Dp \geq d\}, \end{aligned}$$

where $\mathcal{I} \subseteq \{1, \dots, n\}$ is a set of indices, and A, b, D , and d are known.

Our goal is to revisit the main results of §3 and investigate their validity in the present setting. As a general overview, all the critical results can be suitably adapted to the new model: by solving particular MILPs, one can (a) test whether a given x is a PRO solution (and, if not, obtain a PRO solution \bar{x} that Pareto dominates it), (b) generate PRO solutions, and (c) optimize linear objectives over the set of PRO

solutions. The sole results of §3 that no longer hold here are Propositions 1 and 2—in particular, not all PRO solutions can be recovered by maximizing linear objectives of the form $\bar{p}^T x$ where $\bar{p} \in \text{ri}(U)$, and the set X^{PRO} is not necessarily in the boundary of X^{RO} when the two are different. For reasons of brevity, we refrain from reproving the (new) claims, since the arguments exactly parallel their counterparts in §3.

First, note that the representation for the set of robust optimal solutions X^{RO} remains identical to expression (3); i.e.,

$$X^{\text{RO}} = \{x \in X: \exists y \in \mathbb{R}_+^{m_u} \text{ such that } D^T y = x, y^T d \geq z^{\text{RO}}\}.$$

Here, however, the optimal (worst-case) value for the problem, z^{RO} , is obtained by solving an MILP (Bertsimas and Sim 2003), and the set X^{RO} also includes integrality constraints due to X .

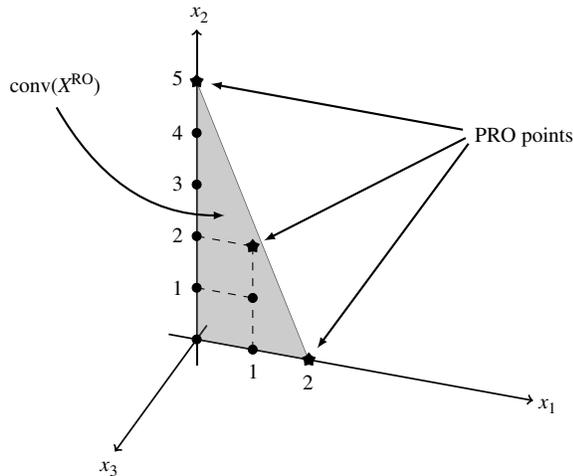
In this context, we can reaffirm our first main result in Theorem 1, which holds true without any modifications. In particular, to test whether a point $x \in X^{\text{RO}}$ is also a PRO, one only needs to solve the optimization problem (7) for an arbitrarily chosen \bar{p} in $\text{ri}(U)$. As before, if the optimal value is exactly zero, then $x \in X^{\text{PRO}}$. Otherwise (i.e., strictly positive optimal value), $x + y^* \in X^{\text{PRO}}$, and it Pareto dominates x . We note that the sole change from the result in §3.1 is that problem (7) is now an MILP instead of an LP. In particular, the computational complexity for finding a PRO solution compared to just finding a solution $x \in X^{\text{RO}}$ is still the same, because both now require solving MILPs. However, the exact computational cost for the former problem may be increased not only because of the requirement of solving additional MILPs but also because it may not preserve some of the structure of the nominal MILP problem.

As before, this theorem readily leads to a simple procedure for finding points $x \in X^{\text{PRO}}$. In particular, a result analogous to Corollary 1 holds: for any $\bar{p} \in \text{ri}(U)$, all the optimal solutions to the MILP maximize $\sum_{x \in X^{\text{RO}}} \bar{p}^T x$ are guaranteed to be PRO points. Furthermore, just as stated in Corollary 2, determining whether $X^{\text{RO}} = X^{\text{PRO}}$ resumes to solving the MILP in (9) and comparing the optimal value with zero.

This result also brings us to the first main point of departure from §3. Recall that by choosing different points $\bar{p} \in \text{ri}(U)$, the union of the sets $\arg \max_{x \in X^{\text{RO}}} \bar{p}^T x$ was guaranteed to contain all PRO solutions for the case without integrality constraints (see Proposition 1 in §3.1). Unfortunately, that is no longer the case here. In particular, there may be points $x \in X^{\text{PRO}}$ that lie in the strict interior of the set $\text{conv}(X^{\text{RO}})$ and hence cannot be recovered by maximizing a linear functional over X^{RO} . The following example presents such an instance.

EXAMPLE 4 (STRICTLY INTERIOR POINT IN X^{RO}). Let $X = \{(x_1, x_2, x_3) \in \mathbb{Z}_+^2 \times \mathbb{R}: \frac{1}{2}x_1 + \frac{1}{5}x_2 \leq 1, x_3 \geq -1,$

Figure 2 Illustration of Sets and Points Considered in Example 4



Notes. The marked points correspond to all the points of X^{RO} ; PRO points are marked as stars. Note that one of the PRO points lies in the relative interior of the convex hull of X^{RO} , which is shaded in gray.

$x_3 \leq 0\}$, and let $U = \text{conv}(\{e_1, e_2, e_3\})$. It can be checked that

$$X^{\text{RO}} = \{x \in X: x_3 = 0\},$$

$$X^{\text{PRO}} = \{[0 \ 2 \ 0]^T, [5 \ 0 \ 0]^T, [1 \ 2 \ 0]^T\}.$$

In particular, $\bar{x} = [1 \ 2 \ 0]^T \in \text{ri}(\text{conv}(X^{\text{RO}}))$, and therefore there is no $p \in \text{ri}(U)$ such that $\bar{x} \in \arg \max_{x \in X^{\text{RO}}} p^T x$. Figure 2 provides a graphical illustration of X^{RO} and X^{PRO} in this case.

Example 4 also proves that there are cases in which $X^{\text{RO}} \neq X^{\text{PRO}}$ and some points in X^{PRO} lie in the strict (relative) interior of X^{RO} such that Proposition 2 also does not hold.

For the same reason, the sampling heuristic suggested in §3.2 for optimizing a linear functional over X^{PRO} may not be very effective: one can certainly still apply it, but since sampling points $\bar{p} \in \text{ri}(U)$ is no longer guaranteed to generate all the points in X^{PRO} , one may be unable to recover the exact (true) optimal value.

However, we note that the problem of optimizing linear objectives over X^{PRO} can nonetheless be dealt with in a scalable fashion. In particular, the main result in Proposition 3 still holds, and, for any $r \in \mathbb{R}^n$, the optimal value in the problem $\max_{x \in X^{\text{PRO}}} r^T x$ can be obtained by solving the MILP in (10). As a result, optimizing over X^{PRO} when some decisions are integral is as easy as finding a single $x \in X^{\text{RO}}$, unlike the setting discussed in §3.2.

5. Numerical Studies

In this section, we evaluate the implications of our findings by performing numerical studies. We focus

on three application areas in which RO has proven to be very powerful. In particular, we study three classical problems from the literature drawn from finance, inventory management, and project management applications.

We generate multiple instances of the problems we consider, using random data, with the purpose of assessing

(a) the frequency at which computed robustly optimal solutions are Pareto dominated by other solutions, and

(b) the performance gain by considering PRO solutions in practice instead of dominated robustly optimal solutions.

For every generated instance, we solve an equivalent problem to problem (9) using the simplex method to identify dominated robustly optimal solutions and associated PRO solutions.

We find that in approximately 8.5% of the instances we generate (across all three problems), there exist robustly optimal solutions that are Pareto dominated by other PRO solutions. Moreover, using the associated PRO solutions in these cases yields a relative performance gain that can be as high as 43% compared with using the dominated robustly optimal solutions.

5.1. Portfolio Optimization

RO has been widely studied and used in financial services applications, particularly in portfolio optimization problems. We refer the reader to Ben-Tal et al. (2000), Goldfarb and Iyengar (2003), Bertsimas and Pachamanova (2008), Natarajan et al. (2008), Calafiore (2007), Ben-Tal et al. (2010), Gabrel et al. (2012) and the book by Fabozzi et al. (2007) for a thorough overview.

In this section, we consider a simple portfolio selection problem that has the form we studied in §2, i.e., a linear program with uncertainty in the objective.

5.1.1. Problem Description. A manager wishes to invest her wealth in $n + 1$ investment opportunities or assets. We denote the return of the i th asset with r_i , $i = 1, \dots, n + 1$. The $(n + 1)$ th asset yields a known, deterministic return of μ_{n+1} , i.e.,

$$r_{n+1} = \mu_{n+1},$$

with no associated risk. On the contrary, all remaining assets are risky. In particular, the return of the i th asset is equal to

$$r_i = \mu_i + \sigma_i \zeta_i, \quad i = 1, \dots, n,$$

where μ_i is the expected return of the i th asset; ζ_i is a random, uncertain shock affecting the return of the i th asset; and σ_i is a volatility parameter. The values of the shocks $\zeta \in \mathbb{R}^n$ are assumed to be bounded between

−1 and 1, and they are assumed to sum up to zero. Uncertainty sets of this style have been introduced and used in Bandi and Bertsimas (2012).

Let $x \in \mathbb{R}^{n+1}$ be the target portfolio composition vector; i.e., x_i is the fraction of the wealth that the manager wishes to invest in the i th asset. We require that (a) no shorting is allowed, i.e., $x \geq 0$; and that (b) the net fraction of the wealth invested in any of the asset groups $\{1, \dots, N/4\}$, $\{N/4, \dots, N/2\}$, $\{N/2, \dots, 3N/4\}$, and $\{3N/4, \dots, N\}$ does not exceed 25%, for diversification purposes. The objective of the manager is to select a portfolio composition so as to maximize its worst-case return.

A formulation of the above RO problem is as follows:

$$\begin{aligned} & \text{maximize} \quad \min_{\zeta \in U} \sum_{i=1}^{n+1} r_i x_i \\ & \text{subject to} \quad r_i = \mu_i + \sigma_i \zeta_i, \quad i = 1, \dots, n, \\ & \quad \quad \quad r_{n+1} = \mu_{n+1}, \\ & \quad \quad \quad \sum_{i=1+(kN)/4}^{((k+1)N)/4} x_i \leq 0.25, \quad k = 0, 1, 2, 3, \\ & \quad \quad \quad \mathbf{1}^T x = 1, \\ & \quad \quad \quad x \geq 0, \end{aligned}$$

with variables $r \in \mathbb{R}^{n+1}$, $x \in \mathbb{R}^{n+1}$, and $U = \{\zeta \in \mathbb{R}^n: -1 \leq \zeta \leq 1, \mathbf{1}^T \zeta = 0\}$.

5.1.2. Data. We consider 10,000 instances of problems of size $n = 8$. The expected returns of the risky assets $\{\mu_i\}_{i=1}^n$ are independently sampled from a uniformly distributed random variable taking values 1%, 1.2%, ..., 3%. The volatility parameters are set to $\sigma_i = 0.2\mu_i + 0.8\xi_i$, for all $i = 1, \dots, n$, where the values $\{\xi_i\}_{i=1}^n$ are independently sampled from the same random variable as the expected returns. The risk-free return is set to $\mu_{n+1} = 1\%$.

5.1.3. Results. The average number of instances in which the solver found a robustly optimal solution x that was not PRO was 31%. For these instances, we obtained a single PRO solution \bar{x} by solving problem (7), for which we recorded two different performance gaps. The first corresponded to the relative improvement that \bar{x} yielded in objective compared with x if a “nominal” scenario \hat{p} materialized, i.e.,

$$\frac{\hat{p}^T(\bar{x} - x)}{\hat{p}^T x},$$

where \hat{p} corresponds to $\zeta = 0$ and the returns equal to their expected values. The second gap corresponded to the maximum relative improvement across all possible scenario realizations, i.e.,

$$\max_{p \in U} \frac{p^T(\bar{x} - x)}{p^T x}.$$

Histograms for both gaps are recorded in Figure 3(a). For the performance gaps under the nominal scenario, the median value recorded was 2.5%, whereas the maximum was as high as 22%. For the maximum performance gaps across all scenarios, the median value recorded was 5%, whereas the maximum was as high as 43%.

5.2. Inventory Management

A different stream of applications in which RO has proven particularly useful has been inventory and supply chain management. For a thorough review and many references, we direct the interested reader to the papers by Ben-Tal et al. (2005), Bertsimas and Thiele (2006), Adida and Perakis (2006), Bienstock and Özbay (2008), Bertsimas et al. (2010, 2011b), and See and Sim (2010); the review papers by Bertsimas et al. (2011a) and Gabrel et al. (2012); and the book by Ben-Tal et al. (2009).

In this section, we consider a single-warehouse multiple-retailer setting, where the manager needs to make stocking decisions in the face of uncertain demand. From an RO perspective, the resulting problem is a two-stage adjustable problem with fixed recourse.

5.2.1. Problem Description. Consider a retail network consisting of a single warehouse and N different retail points, indexed by $i = 1, \dots, N$, where customer demand is realized. For simplicity, we consider the case where a single item is offered for sale.

We have C units of the item available in the warehouse that need to be distributed to the retail points. The i th retail point holds zero initial inventory and is capable of stocking at most c_i units. The transportation costs for distributing or, equivalently, stocking, inventory at the i th retail point is t_i currency units per unit of inventory. Similarly, we assume that there is an operating cost at the i th retail point equal to h_i currency units per unit of inventory.¹¹ The revenues from sales at the i th retail point are r_i currency units per unit of inventory sold.

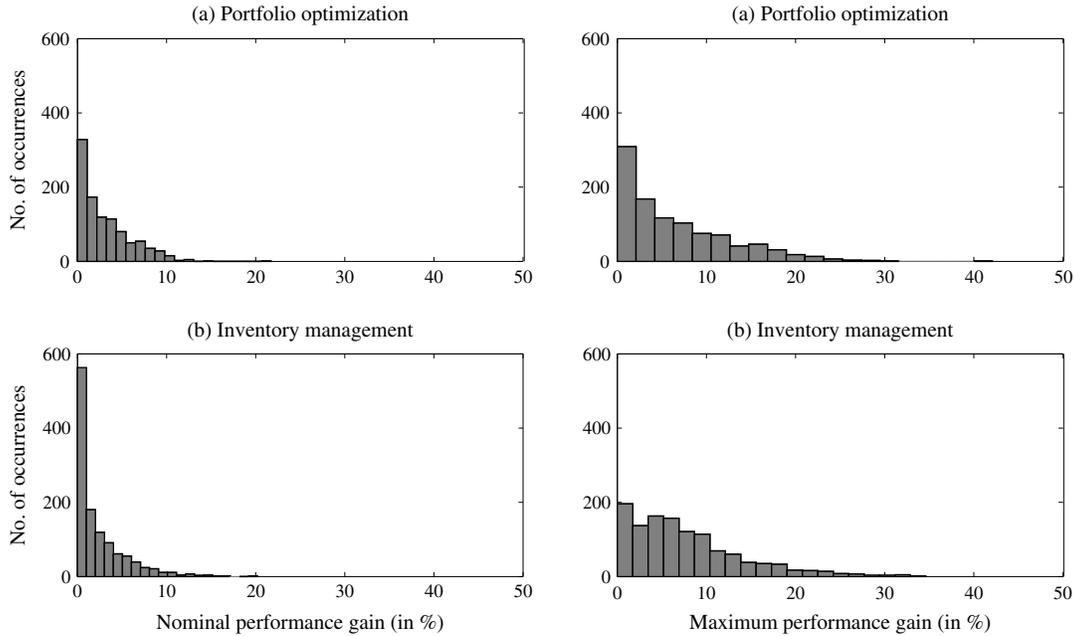
Customer demand at each point is uncertain and is driven by n_f factors that affect the market. In particular, we assume that the demand at the i th point, denoted by d_i , is equal to

$$d_i = d_i^0 + q_i^T z, \quad i = 1, \dots, N,$$

where d_i^0 is the nominal expected demand, $z \in \mathbb{R}^{n_f}$ is a vector of the realized values of the n_f factors, and $q_i \in \mathbb{R}^{n_f}$ is a vector of parameters that measure the exposure of demand d_i to each of the market factors.

¹¹ For simplicity, we assume no fixed transportation or operating costs, although one could include them in a straightforward manner.

Figure 3 Histograms of Nominal (Left) and Maximum (Right) Performance Gains by Considering PRO Solutions in the Numerical Studies of (a) §5.1 and (b) §5.2



The values of the market factors are uncertain. However, we assume that they are bounded, as is their sum, as follows:

$$-b \cdot \mathbf{1} \leq z \leq b \cdot \mathbf{1}, \quad -B \leq \mathbf{1}^T z \leq B,$$

where b and B are known parameters.

Let $s \in \mathbb{R}^N$ be a vector containing the stocking decisions for all points. Similarly, $y \in \mathbb{R}^N$ is a vector with the realized sales at each point. Since customer demand materializes after stocking decisions are made, the sales at the i th point depend on the unknown factors z affecting the demand and are equal to

$$y_i(z) = \min\{s_i, d_i\}, \quad i = 1, \dots, N.$$

The manager makes the stocking decisions so as to maximize worst-case profits, denoted by P . The corresponding RO formulation is as follows:

$$\text{maximize } P \quad (14a)$$

$$\text{subject to } P \leq r^T y(z) - (t+h)^T s \quad \forall z \in U, \quad (14b)$$

$$y(z) \leq s \quad \forall z \in U, \quad (14c)$$

$$y(z) \leq d^0 + Qz \quad \forall z \in U, \quad (14d)$$

$$\mathbf{1}^T s = C, \quad (14e)$$

$$s \leq c, \quad (14f)$$

$$s \geq 0, \quad (14g)$$

with variables $P \in \mathbb{R}$, $y \in \mathbb{R}^N$, and $s \in \mathbb{R}^N$; $U = \{z \in \mathbb{R}^n : -b \cdot \mathbf{1} \leq z \leq b \cdot \mathbf{1}, -B \leq \mathbf{1}^T z \leq B\}$; and $Q = [q_1, \dots, q_N]^T$. Constraint (14b) corresponds to an

upper bound on the worst-case profits, equal to revenues from sales minus transportation and operations costs. According to constraints (14c) and (14d), sales are less than or equal to available stock and realized demand. Constraints (14e) and (14f) correspond to net and individual point capacity constraints, respectively.

Since the sales y depend on the unknown factors z , problem (14) is an adjustable RO problem with fixed recourse. To solve it, we approximate it with its affine adjustable robust counterpart (AARC), a popular and powerful heuristic (see Ben-Tal et al. 2004, 2005; Bertsimas et al. 2010; Iancu et al. 2013; see also the book by Ben-Tal et al. 2009).

The AARC of problem (14) is obtained by substituting the adjustable decisions $y(z)$ with affine functions of z and new auxiliary variables. The resulting formulation is of the same form as problem (11). For our notion of PRO solutions, we require from robustly optimal solutions to maximize slack of constraint (14b). In other words, we employ our methodology from §4.1 using a slack value vector equal to a unit vector, where the component corresponding to constraint (14b) is equal to 1.

Note that slack of constraint (14b) depends on the actual uncertainty realization. If a worst-case demand scenario is realized, the profits on the right-hand side are equal to the worst-case profits P , and the slack is zero. For non-worst-case realizations, however, slack in the constraint is possible, which then translates to additional profits on top of the worst-case value of P (depending on the actual demand realization). PRO

solutions ensure “optimal performance” under any demand realization by maximizing this slack, unlike robustly optimal solutions that ignore it.

5.2.2. Data. We consider 10,000 instances of problems of size $N = 10$. The available inventory is set to $C = 2,000$ units. All other problem data are independently sampled from uniformly distributed random variables. Inventory capacities at individual points $c \in \mathbb{R}^N$ are drawn between 300 and 500. Transportation $t \in \mathbb{R}^N$ and operations cost rates $h \in \mathbb{R}^N$ are drawn between 1 and 3. Sales revenue rates $r \in \mathbb{R}^N$ are drawn between 20 and 40.

Nominal demand values $d^0 \in \mathbb{R}^N$ are drawn between 100 and 200. The number of market factors n_f is drawn from the values 2, 3, and 4. Exposure parameters $Q \in \mathbb{R}^{N \times n_f}$ are drawn between -2 and 2 . Parameters b and B that bound the factor values are set to 5 and 25, respectively.

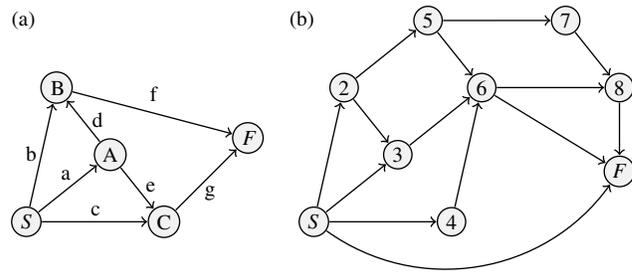
5.2.3. Results. The average number of instances in which a robustly optimal solution was identified that was not a PRO solution was 12%. For these instances, similarly with the study above, we computed the nominal and maximum performance gaps (in terms of actual profits) between the PRO solution and the dominated robustly optimal solution we obtained. In this setting, the nominal scenario corresponds to the case of $z = 0$ and customer demand equal to its expected value. The median performance gap under the nominal scenario recorded was 1.5%, whereas the maximum was as high as 20%. The median of the maximum gaps recorded was 6.5%, whereas the absolute maximum was as high as 35%. Related histograms are depicted in Figure 3(b).

5.3. Project Management

The third application that we consider is focused on robust models for project management. These have been studied in several papers, including Cohen et al. (2007), Goh et al. (2010), Adida and Joshi (2009), and Wiesemann et al. (2012). Here, we consider a model discussed in Ben-Tal et al. (2009), where a manager needs to make resource allocation decisions in the face of uncertain processing times. From a methodological standpoint, the resulting problem is a two-stage adjustable RO problem with fixed recourse.

5.3.1. Problem Description. Consider a project that involves multiple events and activities or tasks to be completed. Projects of this kind are typically analyzed and represented using program evaluation and review technique (PERT) diagrams (Ben-Tal et al. 2009). A *PERT diagram* is a directed, acyclic graph consisting of a set of nodes \mathcal{N} , which correspond to the project events, and a set of edges \mathcal{E} , which correspond to the project activities or tasks.

Figure 4 The PERT Diagrams of the Two Projects Considered in §5.3



Among the nodes there is a start node S and an end node F that correspond to the start and end of the project, respectively. The graph represents logical precedences between the project tasks and events as follows. A task, represented by an edge, starts being processed only after the event that corresponds to the node from which the task originated has occurred. The task is completed after some uncertain processing time. An event occurs only when all the tasks that correspond to all its incoming edges have been completed.

We consider the project of creating a factory studied by Ben-Tal et al. (2009). The project entails the acquirement and delivery of equipment (event A), build of facility 1 (event B), and build of facility 2 (event C). The underlying tasks are as follows: (a) acquiring and delivering of equipment, (b) building facility 1, (c) building facility 2, (d) installing equipment in facility 1, (e) installing equipment in facility 2, (f) training personnel at facility 1, and (g) training personnel at facility 2. The associated PERT diagram is given in Figure 4(a).

To introduce some notation, let task $e \in \mathcal{E}$ originate from node $s(e)$ and terminate at node $f(e)$. Its processing time, denoted by τ_e , is equal to

$$\tau_e = \tau_e^0 + \delta_e, \quad \forall e \in \mathcal{E},$$

where τ_e^0 is the standard processing time of task e and δ_e is an unforeseen delay. We assume that the unforeseen delays of all tasks are bounded, as is their sum, i.e.,

$$\sum_{e \in \mathcal{E}} \delta_e \leq B, \text{ and } 0 \leq \delta_e \leq b, \quad \forall e \in \mathcal{E},$$

where B and b are known parameters.

The manager has the ability of expediting some of the tasks by allocating a scarce resource to them, e.g., by assigning extra workforce or by spending money on purchasing/upgrading task-related equipment. Let z_e denote the amount of the resource allocated to task $e \in \mathcal{E}$. We assume that the resource is divisible.¹² The associated reduction in processing

¹²One could extend our results in case the resource is indivisible by simply enforcing the variables z_e to be discrete.

time of task e is then rz_e , where r is the time reduction rate per unit of resource allocated. The maximum resource amount that can be allocated to any task is equal to c , and the manager has C units of the resource that need to be allocated so as to minimize the worst-case completion time T of the project.

If we denote the time that event $\nu \in \mathcal{N}$ occurs with t_ν , it can be seen that t_ν depends on the uncertain delays δ . More formally, if we let $[S] \stackrel{\text{def}}{=} \emptyset$ and recursively define $[v] \stackrel{\text{def}}{=} \bigcup_{e=(u,v) \in \mathcal{E}} \{[u] \cup \{e\}\}$ as the set of all edges on directed paths from S to v that precede node v , then it can be readily seen that t_ν depends on the delays on the edges in $[v]$, which we denote by $\delta_{[v]}$.

The following formulation captures the manager's problem:

$$\text{minimize } T \quad (15a)$$

subject to

$$t_F(\delta_{[F]}) \leq T, \quad \forall \delta \in U, \quad (15b)$$

$$t_S = 0, \quad (15c)$$

$$t_{f(e)}(\delta_{[f(e)]}) - t_{s(e)}(\delta_{[s(e)]}) \geq \tau_e^0 + \delta_e - rz_e, \\ \forall e \in \mathcal{E}, \forall \delta \in U, \quad (15d)$$

$$\mathbf{1}^T z = C, \quad (15e)$$

$$z \leq c, \quad (15f)$$

$$z \geq 0. \quad (15g)$$

Here, $U = \{\delta \in \mathbb{R}^{|\mathcal{E}|}: 0 \leq \delta \leq b \cdot \mathbf{1}, \mathbf{1}^T \delta \leq B\}$ is the uncertainty set, $T \in \mathbb{R}$ and $z \in \mathbb{R}^{|\mathcal{E}|}$ are nonadjustable decisions,¹³ and the times $t \in \mathbb{R}^{|\mathcal{N}|}$ are adjustable, allowed to depend on the delays realized on edges. Constraints (15b) and (15c) determine the end and start times of the project. Constraint (15d) enforces logical precedence of tasks and events as discussed above. Constraints (15e) and (15f) limit the amount of the resource to be allocated.

Problem (15) is an adjustable RO problem with fixed recourse. To solve it, we approximate it with its AARC similarly to the inventory management problem in the previous section.

For our notion of PRO solutions, we require from robustly optimal solutions to maximize slack of constraint (15b), because it translates to potential reduction of the worst-case value of T (depending on the actual delay realizations). That is, we again employ our methodology from §4.1 using a slack value vector equal to a unit vector, where the component corresponding to constraint (15b) is equal to 1.

¹³ In our model, we assume that the resource allocation decisions z must be made far in advance, before observing the realized delays; in practice, one could also consider adjustable policies. See Ben-Tal et al. (2009) or Wiesemann et al. (2012) for details.

We study separately two PERT graphs, depicted in Figure 4:

(a) the topology from Ben-Tal et al. (2009) we have discussed thus far (panel a), and

(b) a more complicated, randomly generated topology (panel b).

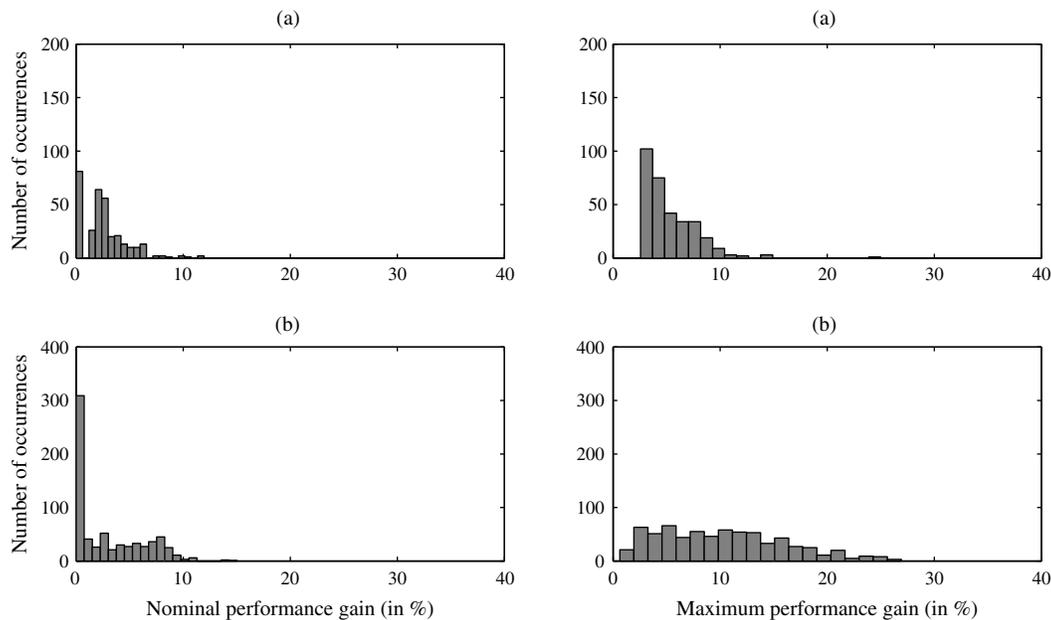
5.3.2. Data. We consider 10,000 instances for each case (a) and (b). The standard processing times $\tau^0 \in \mathbb{R}^{|\mathcal{E}|}$ are independently sampled from a uniformly distributed random variable (a) between 1 and 10, and (b) between 5 and 20. Parameters B and b that bound the unforeseen delays are set to 6 and 2 for case (a), respectively, and 25 and 5 for case (b). The processing time reduction rate r is (a) 1 and (b) 5. The maximum allowable resource at each task c is set to 1 for (a). For case (b), we replace constraint (15f) with $\tau_e^0 \geq rz_e \forall e \in \mathcal{E}$ to ensure nonnegativity of task processing times. The resource amount available C is set to (a) 3 and (b) 10.

5.3.3. Results. The average percentage of instances in which a robustly optimal solution was identified that was not a PRO solution was 3.5% for case (a) and 7% for case (b). Similar to the study above, for these instances we computed the nominal and maximum performance gaps (in terms of actual completion time) between the PRO solution and the dominated robustly optimal solution we obtained. In this setting, the nominal scenario corresponds to the case where δ is the analytic center of U , and tasks are equally delayed. The median performance gap under the nominal scenario recorded was (a) 2.5% and (b) 1.5%, whereas the maximum was as high as (a) 12% and (b) 15%. The median of the maximum performance gaps recorded was (a) 4.5% and (b) 10%, whereas the absolute maximum was as high as (a) 25% and (b) 27%. Related histograms are depicted in Figure 5.

6. Conclusions

In this paper, we adapted the well-known concept of Pareto efficiency in the context of the popular RO methodology. We argued that by focusing exclusively on worst-case outcomes, the classical RO paradigm need not produce solutions that are also Pareto optimal, leading to inefficiencies and suboptimal performance in practice. We provided a basic theoretical characterization of PRO solutions and extended the RO framework by proposing practical methods that verify Pareto optimality, and we generate solutions that are PRO. Critically important, our approach involves solving optimization problems that are of the same complexity as the underlying robust problems; hence, the potential improvements come at essentially limited extra computational cost. Our numerical experiments, drawn from three different application areas (portfolio optimization, inventory management,

Figure 5 Histograms of Nominal (Left) and Maximum (Right) Performance Gains by Considering PRO Solutions in the Numerical Studies of §5.3



and project management), demonstrated that PRO solutions have a significant upside compared with solutions obtained via classical RO methods.

Acknowledgments

The authors gratefully acknowledge the insightful comments received from the editor-in-chief, the associate editor, and four anonymous referees, which have considerably improved the quality of the paper.

References

- Adida E, Joshi P (2009) A robust optimisation approach to project scheduling and resource allocation. *Internat. J. Services, Oper. Informatics* 4(2):169–193.
- Adida E, Perakis G (2006) A robust optimization approach to dynamic pricing and inventory control with no backorders. *Math. Programming* 107(1):97–129.
- Bandi C, Bertsimas D (2012) Tractable stochastic analysis in high dimensions via robust optimization. *Math. Programming* 134(1):23–70.
- Ben-Tal A, Nemirovski A (2002) Robust optimization—Methodology and applications. *Math. Programming* 92(3):453–480.
- Ben-Tal A, Nemirovski A (2007) Selected topics in robust convex optimization. *Math. Programming* 112(1):125–158.
- Ben-Tal A, Bertsimas D, Brown DB (2010) A soft robust model for optimization under ambiguity. *Oper. Res.* 58(4, part 2): 1220–1234.
- Ben-Tal A, El-Ghaoui L, Nemirovski A (2009) *Robust Optimization*, Princeton Series in Applied Mathematics (Princeton University Press, Princeton, NJ).
- Ben-Tal A, Margalit T, Nemirovski A (2000) Robust modeling of multi-stage portfolio problems. Frenk H, Roos K, Terlaky T, Zhang S, eds. *High Performance Optimization* (Kluwer Academic Publishers, Dordrecht, The Netherlands), 303–328.
- Ben-Tal A, Golany B, Nemirovski A, Vial J-P (2005) Retailer-supplier flexible commitments contracts: A robust optimization approach. *Manufacturing Service Oper. Management* 7(3): 248–271.
- Ben-Tal A, Goryashko A, Guslitzer E, Nemirovski A (2004) Adjustable robust solutions of uncertain linear programs. *Math. Programming* 99(2):351–376.
- Bertsimas D, Pachamanova D (2008) Robust multiperiod portfolio management in the presence of transaction costs. *Comput. Oper. Res.* 35(1):3–17.
- Bertsimas D, Sim M (2003) Robust discrete optimization and network flows. *Math. Programming* 98(1–3):49–71.
- Bertsimas D, Thiele A (2006) A robust optimization approach to inventory theory. *Oper. Res.* 54(1):150–168.
- Bertsimas D, Caramanis C, Brown DB (2011a) Theory and applications of robust optimization. *SIAM Rev.* 53(3):464–501.
- Bertsimas D, Farias VF, Trichakis N (2012) On the efficiency-fairness trade-off. *Management Sci.* 58(12):2234–2250.
- Bertsimas D, Iancu DA, Parrilo PA (2010) Optimality of affine policies in multistage robust optimization. *Math. Oper. Res.* 35(2):363–394.
- Bertsimas D, Iancu DA, Parrilo PA (2011b) A hierarchy of near-optimal policies for multistage adaptive optimization. *IEEE Trans. Automatic Control* 56(12):2809–2824.
- Bertsimas D, Litvinov E, Sun X, Zhao J, Zheng T (2013) Adaptive robust optimization for the security constrained unit commitment problem. *IEEE Trans. Power Systems* 28(1):52–63.
- Bienstock D, Özbay N (2008) Computing robust basestock levels. *Discrete Optim.* 5(2):389–414.
- Boyd S, Vandenberghe L (2004) *Convex Optimization* (Cambridge University Press, Cambridge, UK).
- Calafiore G (2007) Ambiguous risk measures and optimal robust portfolios. *SIAM J. Optim.* 18(3):853–877.
- Chan TC, Bortfeld T, Tsitsiklis JN (2006) A robust approach to IMRT optimization. *Physics Medicine Biol.* 51(10):2567–2583.
- Chen W, Unkelbach J, Trofimov A, Madden T, Kooy H, Bortfeld T, Craft D (2012) Including robustness in multi-criteria optimization for intensity-modulated proton therapy. *Physics Medicine Biol.* 57(3):591–608.
- Cohen I, Golany B, Shtub A (2007) The stochastic time-cost trade-off problem: A robust optimization approach. *Networks* 49(2): 175–188.
- Cristobal MA, Parks GT, Clarkson PJ (2006) Finding robust solutions to multi-objective optimisation problems using polynomial chaos. *Proc. 6th ASMO UK/ISSMO Conf. Engrg. Design Optim.* CD-ROM. ASMO UK, Oxford, UK.

- Deb K, Gupta H (2005) Searching for robust Pareto-optimal solutions in multi-objective optimization. Coello Coello CA, Hernández Aguirre A, Zitzler E, eds. *Evolutionary Multi-Criterion Optimization*, Lecture Notes in Computer Science, Vol. 3410 (Springer, Berlin), 150–164.
- El-Ghaoui L, Oustry F, Lebret H (1998) Robust solutions to uncertain semidefinite programs. *SIAM J. Optim.* 9(1):33–52.
- Fabozzi F, Kolm P, Pachamanova D, Focardi S (2007) *Robust Portfolio Optimization and Management*. Frank J. Fabozzi Series (John Wiley & Sons, Hoboken, NJ).
- Fischetti M, Monaci M (2009) Light robustness. Ahuja R, Möhring R, Zorliagis C, eds. *Robust and Online Large-Scale Optimization*, Lecture Notes in Computer Science, Vol. 5868 (Springer, Berlin), 61–84.
- Gabrel V, Murat C, Thiele A (2012) Recent advances in robust optimization and robustness: An overview. Technical report, LAMSADE, Université Paris-Dauphine, Paris.
- Gabrel V, Murat C, Wu L (2013) New models for the robust shortest path problem: Complexity, resolution and generalization. *Ann. Oper. Res.* 207(1):97–120.
- Goh J, Hall N, Sim M (2010) Robust optimization strategies for total cost control in project management. Working paper, Stanford University, Stanford, CA.
- Goldfarb D, Iyengar G (2003) Robust portfolio selection problems. *Math. Oper. Res.* 28(1):1–38.
- Gorissen BL, den Hertog D (2011) Approximating the Pareto set of multiobjective linear programs via robust optimization. CentER Discussion Paper 2012-031, Center for Economic Research, Tilburg University, Tilburg, The Netherlands.
- Iancu DA, Sharma M, Sviridenko M (2013) Supermodularity and affine policies in dynamic robust optimization. *Oper. Res.* 61(4):941–956.
- Joshi S, Boyd S (2009) An efficient method for large-scale slack allocation. *Engrg. Optim.* 41(12):1163–1176.
- Kalai R, Lamboray R, Vanderpooten D (2012) Lexicographic α -robustness: An alternative to min–max criteria. *Eur. J. Oper. Res.* 220(3):722–728.
- Leung JY-T, Pinedo M, Wan G (2010) Competitive two-agent scheduling and its applications. *Oper. Res.* 58(2):458–469.
- Lim AEB, Shanthikumar JG (2007) Relative entropy, exponential utility, and robust dynamic pricing. *Oper. Res.* 55(2):198–214.
- Luo B, Zheng J (2008) A new methodology for searching robust Pareto optimal solutions with MOEAs. *Proc. IEEE World Congress Comput. Intelligence* (IEEE, Piscataway, NJ), 580–586.
- Luo Z-Q, Pang J-S, Ralph D (1996) *Mathematical Programs with Equilibrium Constraints* (Cambridge University Press, Cambridge, UK).
- Natarajan K, Pachamanova D, Sim M (2008) Incorporating asymmetric distributional information in robust value-at-risk optimization. *Management Sci.* 54(3):573–585.
- Ogryczak W (1997) On the lexicographic minimax approach to location problems. *Eur. J. Oper. Res.* 100(3):566–585.
- Ono S, Nakayama S (2009) Multi-objective particle swarm optimization for robust optimization and its hybridization with gradient search. *IEEE Congress Evolutionary Comput.* (IEEE, Piscataway, NJ), 1629–1636.
- Perakis G, Roels G (2010) Robust controls for network revenue management. *Manufacturing Service Oper. Management* 12(1): 56–76.
- Rockafellar T (1970) *Convex Analysis* (Princeton University Press, Princeton, NJ).
- Roy B (2010) Robustness in operational research and decision aiding: A multi-faceted issue. *Eur. J. Oper. Res.* 200(3):629–638.
- Rusmevichientong P, Topaloglu H (2012) Robust assortment optimization in revenue management under the multinomial logit choice model. *Oper. Res.* 60(4):865–882.
- Savage LJ (1972) *The Foundations of Statistics*, 2nd ed. (Dover Publications, Mineola, NY).
- Schrijver A (2000) *Theory of Linear and Integer Programming*, 2nd ed. (John Wiley & Sons, Chichester, UK).
- See C-T, Sim M (2010) Robust approximation to multiperiod inventory management. *Oper. Res.* 58(3):583–594.
- Soyster AL (1973) Convex programming with set-inclusive constraints and applications to inexact linear programming. *Oper. Res.* 21(5):1154–1157.
- Sturm J, Polik I, Terlaky T (2006) SeDuMi 1.3. <http://sedumi.ie.lehigh.edu/>.
- Suh M, Lee T (2001) Robust optimization method for the economic term in chemical process design and planning. *Indust. Engrg. Chemistry Res.* 40(25):5950–5959.
- Wiesemann W, Kuhn T, Rustem B (2012) Robust resource allocations in temporal networks. *Math. Programming* 135(1–2): 437–471.
- Ye Y (1992) On the finite convergence of interior-point algorithms for linear programming. *Math. Programming* 57(1–3):325–335.
- Young HP (1995) *Equity: In Theory and Practice* (Princeton University Press, Princeton, NJ).