

Inverse Optimization with Discrete Decisions

Jerry Chua

Institute of Operations Research and Analytics, National University of Singapore
Institute for Infocomm Research (I2R), Agency for Science, Technology and Research (A*STAR)
jerry.chua@u.nus.edu

Ian Yihang Zhu

Department of Analytics and Operations, NUS Business School, National University of Singapore
ianyzhu@nus.edu.sg

Inverse optimization (IO) has emerged as a powerful framework for analyzing prescriptive model parameters that rationalize observed or prescribed decisions. Despite the prevalence of discrete decision-making models, existing work has primarily focused on continuous and convex models, for which the corresponding IO problems are far easier to solve. This paper makes three contributions that broaden the foundations and applications of inverse optimization in discrete decision-making. First, we propose a new approach for approximating the inverse-feasible region of solutions to discrete optimization problems by modifying their linear relaxations. Second, we integrate these modified relaxations with existing methods to develop a unified framework for solving IO problems. Finally, through the lens of inverse optimization, we provide a fresh perspective on the sensitivity analysis of prescriptive solutions to discrete optimization problems, and demonstrate the computational advantages of our models and algorithms.

Key words: inverse optimization, discrete optimization, prescriptive analytics, linear relaxations, decomposition methods, sensitivity and post-optimality analysis

History: November 10, 2025.

1. Introduction

In recent years, inverse optimization (IO) has attracted significant attention in the operations research and prescriptive analytics communities (Chan et al. 2025). It has emerged as a powerful framework for analyzing prescriptive model parameters that explain and rationalize observed or prescribed decisions. While several modeling paradigms have been proposed, most existing work focuses on IO problems defined in their canonical form: given a feasible solution to a nominal optimization problem – known as the *forward problem* – the goal of the IO problem is to determine the objective coefficients of this forward problem under which the given solution is rendered optimal. When multiple objective coefficient vectors satisfy this condition, we seek the vector that optimizes a desired criterion while respecting any context-specific constraints on the coefficients.

IO problems defined in this canonical form can be found in a variety of contexts. For example, they appear widely in bilevel optimization and incentive design problems, where they are used to compute prices or incentives that render prescribed decisions optimal for decision-making agents

(e.g., Agarwal and Ergun 2010, Liu et al. 2024). Similarly, in data-driven and contextual optimization problems, they can be applied to infer latent utility coefficients that rationalize certain decisions (e.g., Lin et al. 2024), or to train machine learning models to make decision-focused predictions (e.g., Tang and Khalil 2024). These IO models have been applied across diverse sectors spanning energy, transportation, and healthcare systems; Chan et al. (2025) provides a comprehensive overview of such applications.

The defining feature of these IO problems is the *inverse-feasible region* (IFR), which describes the complete set of objective coefficients that render a given solution optimal to the forward problem (Tavashoghlu et al. 2018). By design, canonical IO problems are restricted to generate *inverse-feasible* vectors – that is, vectors lying in the inverse-feasible region – and the tractability of solving these problems thus depends on how easily this region can be represented. From this perspective, all IO solution methods can also be viewed through the way they represent or explore this region.

When the forward problem is a linear program, the inverse-feasible region can be represented using standard optimality conditions. In particular, objective coefficients that make a given solution optimal to the forward problem are precisely those that satisfy strong duality or complementary slackness conditions (Ahuja and Orlin 2001). In this context, IO problems are straightforward to solve, and the vast majority of IO applications have focused exclusively on forward problems that are linear programs (Chan et al. 2025).

When the forward problem contains discrete decision variables, the inverse-feasible region becomes more difficult to handle, and the corresponding IO problems are significantly harder to solve. While exact reformulations of this region have been proposed – for example, through a super exponential set of subadditive functions (Schaefer 2009, Lamperski and Schaefer 2015) – they are largely conceptual in nature and difficult to implement in practice. Instead, cutting plane algorithms stand as the predominant approach for solving these IO problems, where conditions defining the inverse-feasible region are relaxed and cuts are iteratively generated to exclude candidate vectors that are not inverse-feasible (Wang 2009, Bodur et al. 2022).

While cutting plane algorithms have shown some promise, they suffer from a fundamental limitation: inverse-feasible vectors are, by definition, neither identified nor improved upon until the algorithm converges to an optimal solution to the IO problem. In practice, this means that regardless of whether the solution process takes minutes, hours, or days, nothing informative is generated until the process terminates. This “all-or-nothing” approach can be highly limiting, most notably in settings where exact solution times may far exceed practical or reasonable time limits.

1.1. Summary of contributions

In this paper, we focus on solving canonical IO problems when the forward problem contains discrete decision variables. In this context, we make three major contributions. First, we propose a

novel approach for approximating the inverse-feasible region by leveraging the linear relaxation of the forward problem. Second, we use these approximations to develop a generalized framework for solving IO problems, designed to iteratively generate and improve inverse-feasible vectors. Finally, we introduce novel applications of inverse optimization in the sensitivity and post-optimality analysis of discrete optimization problems, through which we derive insights into our methods and algorithms. A more detailed summary of each contribution is given below.

I. Tractable approximations of the inverse-feasible region. Mathematically, our approach is motivated by a rather simple question: for a discrete optimization problem, how can we best leverage its linear relaxation to approximate the inverse-feasible region? While linear relaxations have long served as a cornerstone of solution methods for discrete optimization, their role in inverse optimization has received surprisingly little attention. We demonstrate that typical linear relaxations may generate poor approximations of the inverse-feasible region, but they can be restricted in a way that makes them significantly more effective for approximating this region. Specifically, we introduce the concept of *restricted linear relaxations* (RLRs), which are defined by adding a family of inequalities that are not valid for the linear relaxation. We demonstrate the advantages and versatility of RLRs through a number of theoretical results and practical examples.

II. Generalized solution methods for inverse optimization. We leverage RLRs to introduce a generalized framework for solving inverse discrete optimization problems. Specifically, we integrate RLRs into existing cutting plane approaches, which enables inverse-feasible solutions to be generated and refined iteratively. In this integrated approach, the RLR provides a unique and versatile mechanism for restricting the search space of solutions, allowing users to control the trade-off between pursuing small versus large solution improvements, and to specify termination conditions based on optimality gaps. Conceptually, this generalization can be viewed as a more complete primal-dual solution framework for inverse optimization.

III. New applications and computational insights. We broaden the scope of inverse optimization by demonstrating that a large family of *sensitivity and post-optimality analysis* (SPA) problems can be posed and solved as canonical IO problems. To the best of our knowledge, we are the first to bridge these two areas of research. By forming this connection, we show that our solution methods provide a unified framework for solving SPA problems in discrete optimization, covering both classic and modern SPA problems that we introduce for multi-objective and data-driven optimization. Using one such problem, we derive insights into the characteristics and performance of our solution methods. Our results demonstrate that under strict time limits, existing approaches can easily become irrelevant, while our generalized approach can still generate near-optimal solutions. We also investigate trade-offs between obtaining inverse-feasible solutions quickly and dedicating more time to achieving significantly better solutions.

The rest of our paper is organized as follows. Section 2 provides an overview of the relevant literature. In Section 3, we define the inverse optimization problem and present preliminary insights that lay the foundations for developments in subsequent sections. Section 4 introduces restricted linear relaxations and examines their use in approximating the inverse-feasible region. Section 5 outlines the generalized framework for solving IO problems by integrating RLRs with cutting plane approaches. Section 6 focuses on applications and numerical experiments. Finally, Section 7 discusses directions for future research. All proofs can be found in the Electronic Companion.

1.2. Notation

We highlight common notation used in this paper. We use boldface uppercase letters for matrices (e.g., \mathbf{A}), boldface lowercase letters for vectors (e.g., \mathbf{b}), and calligraphic letters for sets (e.g., \mathcal{X}). The vector of all zeros (ones) is written as $\mathbf{0}$ ($\mathbf{1}$), and the i -th basis vector is given as \mathbf{e}_i . The sets \mathbb{R}_+ and \mathbb{R}_- represent the positive and negative orthant, respectively. For a matrix \mathbf{A} , the i -th row vector is denoted as $\mathbf{a}_i \in \mathcal{A}$ and the i -th column vector is denoted as $\mathbf{A}_i \in \mathcal{A}$. For any set of vectors $\mathbf{a}_1, \dots, \mathbf{a}_m$, its conic hull is given as $\text{cone}(\mathbf{a}_1, \dots, \mathbf{a}_r) = \{\mathbf{x} \in \mathbb{R}^n \mid \mathbf{x} = \sum_{i=1}^r \lambda_i \mathbf{a}_i, \lambda_i \in \mathbb{R}_+\}$. For a feasible region \mathcal{X} , we denote \mathcal{X}^{LR} as the linear relaxation of \mathcal{X} , $\text{conv}(\mathcal{X})$ as its convex hull, and $\text{ext}(\mathcal{X})$ as the extreme points of this convex hull. Finally, for a polyhedral cone \mathcal{K} , we denote $\dim(\mathcal{K})$ as the dimensionality of the cone, and $\text{int}(\mathcal{K})$ to be the interior of this cone, if it exists.

2. Literature Review

We first review the relevant literature on inverse optimization in Section 2.1, followed by the literature on sensitivity and post-optimality analysis in Section 2.2. These two areas have, to date, been largely examined independently. In both subsections, we mainly focus on work pertaining to discrete optimization. In Section 2.3, we briefly outline additional problems involving the inverse-feasible region that may be of broader interest.

2.1. Inverse optimization

The IO literature has developed along two main streams of research that differ in whether the inverse-feasibility condition is imposed, that is, whether the output of the IO problem must render a given solution optimal to the forward problem (Chan et al. 2025). The first stream, which we refer to as canonical IO problems, enforces this condition, while the second stream, often termed as data-driven IO problems, relaxes this requirement. These two paradigms differ primarily in their intended applications. In this paper, we focus on canonical IO problems and review this body of work, while a brief discussion of data-driven IO is provided at the end of the subsection.

To date, most of the IO literature focuses on forward problems that are linear programs. For example, nascent literature focused on special cases of inverse linear programs for which one can

derive analytical solutions based on the value of the optimal dual variables (e.g., [Ahuja and Orlin 2001](#), [Heuberger 2004](#)). Similarly, many applications of inverse optimization focus on problems where decision-making agents are assumed to be solving continuous and convex optimization problems (e.g., [Marcotte et al. 2009](#), [Agarwal and Ergun 2010](#)). In such cases, strong duality and complementary slackness can be exploited to derive tractable reformulations of these IO problems.

Inverse discrete optimization problems, which have been of more recent interest, are significantly harder to solve due to the difficulty of reformulating the inverse-feasible region ([Schaefer 2009](#), [Lamperski and Schaefer 2015](#), [Bulut and Ralphs 2021](#)). As mentioned in the introduction, cutting plane algorithms serve as the predominant approach for solving these problems in practice (e.g., [Wang 2009](#), [Bodur et al. 2022](#)), and we will show that they iteratively refine an outer approximation of the inverse-feasible region. While this approach has shown some promise, it also possesses inherent limitations. Specifically, because the approach successively refines outer approximations of the inverse-feasible region, no inverse-feasible solutions are found until the algorithm converges to the optimal solution. This results in an “all-or-nothing” solution process that is entirely uninformative until the algorithm converges, a process which often requires a significant amount of time. The only paper that has aimed to address this challenge is [Duan and Wang \(2011\)](#), which proposes a regularization-based heuristic designed to run in parallel with the cutting plane algorithm for solving a particular type of inverse optimization problem.

In this paper, we take a holistic approach to analyzing and solving inverse discrete optimization problems. We first propose a simple, versatile, and tractable method for approximating the inverse-feasible region by leveraging the linear relaxation of the underlying discrete optimization problem. We then integrate these approximation methods into a unified solution framework for inverse optimization that addresses the main limitations discussed above.

IO problems in their canonical form arise in diverse domains including incentive design, network pricing, decision-focused learning, cooperative games, and bilevel optimization; we refer to the survey paper of [Chan et al. \(2025\)](#) for a comprehensive overview. While most of these applications assume that agents solve convex optimization problems, they can be naturally extended to richer and more realistic environments where there are discrete decisions. As we will discuss in [Section 2.2](#), our paper also broadens the scope of applications by showing that a large family of sensitivity and post-optimality analysis (SPA) problems can be formulated and solved as IO problems.

Finally, as noted earlier, there is also a growing body of literature on data-driven IO models (e.g., [Aswani et al. 2018](#), [Mohajerin Esfahani et al. 2018](#), [Shahmoradi and Lee 2022](#), [Zattoni Scroccaro et al. 2025](#)). These models are designed primarily for statistical and supervised learning, where the goal is to learn unknown parameters of a forward problem by minimizing suboptimality-based loss functions ([Chan et al. 2025](#)). Although this use case differs from the primary applications of

canonical IO models, the latter may still be relevant, both for the sensitivity analysis of estimated parameters, and in situations where the canonical IO formulation appears as a subproblem within the parameter estimation process (e.g., [Chan et al. 2024](#)).

2.2. Sensitivity and post-optimality analysis

Inverse optimization is closely related to the domain of sensitivity and post-optimality analysis (SPA), although the two areas have evolved largely independently. For example, in the context of SPA, the inverse-feasible region is sometimes referred to as the sensitivity or stability region (e.g., [Kılınç-Karzan et al. 2009](#), [Andersen et al. 2023](#)). For a given solution to nominal optimization problem, many SPA problems aim to determine the ranges over which model parameters can vary without affecting the optimality of the given solution. Among these, the most widely studied is *single-coefficient* sensitivity analysis problem, which seeks to identify the largest interval within which a single objective coefficient can vary while all others are held fixed. The inverse-feasible region by definition contains the cartesian product of all such intervals.

SPA problems for linear programs are straightforward to solve, and the common example of the single-coefficient sensitivity analysis problem is widely applied in practice and even taught in introductory optimization courses ([Bertsimas and Tsitsiklis 1997](#)). In contrast, SPA problems are considerably more challenging in the context of discrete optimization, where existing methods can often be highly problem-specific and analytically demanding while lacking in generalizability. For example, to solve the single-coefficient sensitivity analysis problem, early methods focus on techniques that leverage bounds on the objective value in each node of a pure branch-and-bound tree that we assume is used to generate a solution to the discrete optimization problem ([Roodman 1974](#), [Geoffrion and Nauss 1977](#), [Jenkins 1982](#), [Schrage and Wolsey 1985](#)). Conversely, [Klein and Holm \(1979\)](#) shows that if the discrete optimization problem is instead solved using a pure cutting plane method, then the cuts that are binding at the optimal solution generates a valid region of objective coefficients for which this solution remains optimal. More recently, [Kılınç-Karzan et al. \(2009\)](#) focuses on computing sensitivity ranges for the objective coefficients of binary variables in a general discrete optimization problem. They show that a valid sensitivity range can be computed by taking the difference in objective values when the corresponding binary variable is forced to take the opposite value of the one that appears in the optimal solution. [Andersen et al. \(2023\)](#) proposes an alternative approach that reframes the sensitivity analysis problem as a bi-objective optimization problem, which is then solved using custom multi-objective optimization methods. Finally, attempts to approximate the inverse-feasible region have also been proposed for specific problem structures. For example, [Niendorf et al. \(2015\)](#) and [Niendorf and Girard \(2017\)](#) examine the inverse-feasible region for the traveling salesman problem and propose approximations based on graph-theoretic properties specific to their problem setting.

In this paper, we will show that a wide range of classic and modern SPA problems can be formulated as a canonical IO problem. This provides a unified methodological framework for analyzing the solutions to prescriptive models that contain discrete decision variables.

2.3. Additional applications of the IFR

The inverse-feasible region is also emerging in new applications that lie beyond the scope of the IO models considered in this paper. We include them below for broader interest, and as potential avenues for future research (see Section 7). For example, Wei and Zhang (2024) study conditions under which adjustable robust optimization problems yields the same objective value as its static counterpart, and show that the conditions depend on the IFR of a point in the polyhedral uncertainty set. Besbes et al. (2025) study contextual optimization problems where feedback is given in terms of optimal actions. Their optimal policy is based on approximations of the IFR of these actions. Chan et al. (2024) propose robust optimization models for recommending decisions using observations of past decisions. Their approach calibrates an uncertainty set on objective coefficients by requiring that this set intersects with IFRs associated with past decisions. Tang and Khalil (2024) and Berden et al. (2025) propose decision-focused methods for training models that predict the objective coefficients of linear optimization problems. Their approach aims to minimize deviations of these predictions from the IFR of optimal solutions. In the same context, Zhou et al. (2025) quantifies the risk associated with a prescriptive decision based on its inverse-feasible region.

3. Problem Definition and Preliminary Insights

We first begin by motivating and defining an inverse optimization problem in Section 3.1. We then outline preliminary ideas for approximating and solving these problems in Section 3.2. Section 3.3 provides a short summary of the main insights and presents a roadmap for the rest of the paper.

3.1. Problem definition and motivation

Consider a general discrete optimization problem of the form

$$\begin{aligned} \text{FO}(\boldsymbol{\theta}, \mathcal{X}) &:= \max_{\mathbf{x}} \quad \boldsymbol{\theta}^\top \mathbf{x} \\ \text{s.t.} \quad &\mathbf{x} \in \mathcal{X} := \left\{ \mathbf{x} \mid \mathbf{A}\mathbf{x} \leq \mathbf{b}, \mathbf{x} \in \mathbb{R}^\ell \times \mathbb{Z}^{n-\ell} \right\}. \end{aligned} \tag{FO}$$

Throughout this paper, we refer to this nominal problem as the forward problem (FO), and to the set \mathcal{X} as the forward-feasible region. We assume \mathcal{X} is bounded. Then, given a solution $\hat{\mathbf{x}} \in \mathcal{X}$, the inverse-feasible region of $(\hat{\mathbf{x}}, \mathcal{X})$ is defined as the set of all possible vectors $\boldsymbol{\theta}$ under which $\hat{\mathbf{x}}$ is optimal in the forward problem, i.e.,

$$\mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X}) := \left\{ \boldsymbol{\theta} \mid \hat{\mathbf{x}} \in \arg \max \text{FO}(\boldsymbol{\theta}, \mathcal{X}) \right\}. \tag{RINV}$$

The following example provides a simple illustration of the forward and inverse-feasible region of a discrete optimization problem.

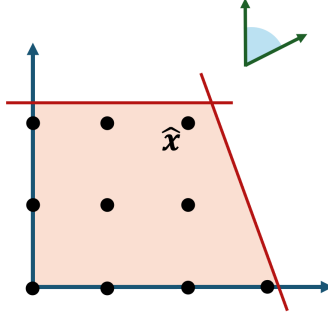


Figure 1 A forward-feasible solution $\hat{x} \in \mathcal{X}$ and its corresponding inverse-feasible region $\mathcal{R}_{\text{inv}}(\hat{x}, \mathcal{X})$.

Example 1. Figure 1 depicts a forward-feasible region $\mathcal{X} := \{x \in \mathbb{Z}^2 \mid x_1 + 0.4x_2 \leq 3.1, x_2 \leq 2.2\}$ and a feasible solution $\hat{x} = (2, 2)$. The inverse-feasible region $\mathcal{R}_{\text{inv}}(\hat{x}, \mathcal{X})$ is illustrated by the blue cone in the top right corner, where any vector θ drawn from this cone will render \hat{x} an optimal solution to the forward problem $\text{FO}(\theta, \mathcal{X})$. In this figure, the shaded polyhedron represents the linear relaxation of \mathcal{X} , written as $\mathcal{X}^{\text{LR}} = \{x \in \mathbb{R}^2 \mid x_1 + 0.4x_2 \leq 3.1, x_2 \leq 2.2\}$.

We now define a general inverse optimization problem of the form

$$\begin{aligned} \text{IO}(\hat{x}, \mathcal{X}) &:= \min_{\theta, s} f(s) \\ \text{s.t. } &\theta \in \mathcal{R}_{\text{inv}}(\hat{x}, \mathcal{X}) \\ &\theta \in \Theta(s) \\ &s \in \mathcal{S}. \end{aligned} \tag{IO}$$

In (IO), the main decision variables are θ , which correspond to the objective vector of the forward problem, and the first constraint ensures that θ must lie in the inverse-feasible region $\mathcal{R}_{\text{inv}}(\hat{x}, \mathcal{X})$. Problem (IO) also contains auxiliary variables $s \in \mathcal{S}$, where $\Theta(s)$ defines a linear mapping $\mathcal{S} \rightarrow \mathbb{R}^n$ from s to θ , while $f(s)$ denotes the objective function to be optimized. Although these auxiliary variables are not required in all applications, they are included for generality and are particularly useful in the problems that we will examine; an example will be given shortly. Unless stated otherwise, we assume that $f(s)$ is a convex function and that \mathcal{S} is a convex set.

Remark 1. The formulation above generalizes inverse optimization models without auxiliary variables, which can be expressed using $\Theta(s) := \mathbf{I}s$, where $s \in \mathbb{R}^n$ and \mathbf{I} denotes the identity matrix.

As we have discussed in Sections 1 and 2, problems in the form (IO) arise across diverse applications. In this paper, we broaden the scope of these applications by showing that (IO) can model and solve a wide range of sensitivity and post-optimality analysis (SPA) problems in discrete optimization. Specifically, SPA problems commonly aim to derive conditions under which a solution $\hat{x} \in \mathcal{X}$ remains optimal when nominal model parameters are perturbed. The most commonly studied case is the single-coefficient sensitivity analysis problem, which seeks to determine the largest interval

within which a single objective coefficient θ_i can deviate from its nominal value while the solution remains optimal to the forward problem (see Section 2.2). This is in fact an inverse optimization problem. Specifically, let $\hat{\mathbf{x}}$ denote an optimal solution to $\text{FO}(\hat{\boldsymbol{\theta}}, \mathcal{X})$ under a nominal cost vector $\hat{\boldsymbol{\theta}}$. For each coefficient $\hat{\theta}_i$, we can compute a maximal range $[\hat{\theta}_i - l_i, \hat{\theta}_i + u_i]$ for which $\hat{\mathbf{x}} \in \mathcal{X}$ remains optimal by solving a pair of inverse optimization problems, i.e.,

$$l_i = \min_{\boldsymbol{\theta}, s} \{s \mid \boldsymbol{\theta} = \hat{\boldsymbol{\theta}} + s\mathbf{e}_i, \boldsymbol{\theta} \in \mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})\} \quad (1a)$$

$$u_i = \max_{\boldsymbol{\theta}, s} \{s \mid \boldsymbol{\theta} = \hat{\boldsymbol{\theta}} + s\mathbf{e}_i, \boldsymbol{\theta} \in \mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})\}, \quad (1b)$$

where \mathbf{e}_i denotes a unit vector where all elements $j \neq i$ are equal to zero.

In Section 6, we will show that many other SPA problems can be specified using model (IO). To the best of our knowledge, we are the first to bridge these two areas of research. We will demonstrate that this connection facilitates a unified, IO-based framework for solving a wide range of SPA problems, including new ones designed for modern decision-making paradigms.

3.2. Preliminary observations on the IFR

As we have discussed in the introduction, the defining feature of inverse optimization problems is the inverse-feasible region $\mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$. To facilitate the upcoming discussion, we first provide some preliminary observations and insights into approximations of this region. We first begin with a simple but important observation, which is that $\mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$ is a polyhedral cone that can be represented through a set of supporting hyperplanes.

Observation 1. *The inverse-feasible region $\mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$ is equivalent to the polyhedral cone*

$$\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}) := \{\boldsymbol{\theta} \mid \boldsymbol{\theta}^\top (\hat{\mathbf{x}} - \mathbf{x}) \geq 0, \forall \mathbf{x} \in \mathcal{X}\}. \quad (2)$$

Equation (2) reformulates the optimality condition in (R_{INV}) explicitly using a set of hyperplanes which ensure that $\hat{\mathbf{x}}$ attains an objective value no smaller than that of any other solution $\mathbf{x} \in \mathcal{X}$. In general, this exact representation cannot be used directly, as it is based on an exponentially large or semi-infinite set of hyperplanes, each defined by a feasible solution of \mathcal{X} that must first be generated. Nonetheless, the characterization provides a valuable perspective for approximating the inverse-feasible region, and one that sets the stage for the rest of the paper. Specifically, we can approximate $\mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$ by replacing the \mathcal{X} in $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X})$ with a surrogate set $\tilde{\mathcal{X}}$.

Remark 2. *From this point onward, we use the notation $\mathcal{K}(\hat{\mathbf{x}}, \cdot)$ to refer to approximations of $\mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$. This choice is deliberate: for a given surrogate $\tilde{\mathcal{X}}$, the sets $\mathcal{K}(\hat{\mathbf{x}}, \tilde{\mathcal{X}})$ and $\mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \tilde{\mathcal{X}})$ coincide when $\hat{\mathbf{x}} \in \tilde{\mathcal{X}}$, but differ when $\hat{\mathbf{x}} \notin \tilde{\mathcal{X}}$, since $\mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \tilde{\mathcal{X}}) = \emptyset$ whereas $\mathcal{K}(\hat{\mathbf{x}}, \tilde{\mathcal{X}}) \neq \emptyset$.*

3.2.1. Approximation I: subsets of \mathcal{X} and cutting plane algorithms. The most natural example of a surrogate set is $\tilde{\mathcal{X}} \subseteq \mathcal{X}$, where $\tilde{\mathcal{X}}$ defines any subset of solutions known to lie in \mathcal{X} . By definition, $\mathcal{K}(\hat{\mathbf{x}}, \tilde{\mathcal{X}})$ would then be an outer approximation of $\mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$, as $\mathcal{K}(\hat{\mathbf{x}}, \tilde{\mathcal{X}})$ would be defined by a subset of the supporting hyperplanes from (2). Similarly, for any $\tilde{\mathcal{X}}^2 \subseteq \tilde{\mathcal{X}}^1 \subseteq \mathcal{X}$, it must be true that $\mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X}) \subseteq \mathcal{K}(\hat{\mathbf{x}}, \tilde{\mathcal{X}}^1) \subseteq \mathcal{K}(\hat{\mathbf{x}}, \tilde{\mathcal{X}}^2)$. Practically, this implies that enumerating more feasible solutions in \mathcal{X} yields progressively tighter outer approximations of $\mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$.

Cutting plane (CP) approaches for inverse optimization operate precisely by enumerating solutions in \mathcal{X} . Specifically, the standard CP approach alternates between a master problem that solves a relaxed version of (IO) where $\mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$ is replaced with $\mathcal{K}(\hat{\mathbf{x}}, \tilde{\mathcal{X}})$ for some $\tilde{\mathcal{X}} \subset \mathcal{X}$, i.e.,

$$\begin{aligned} \min_{\boldsymbol{\theta}, \mathbf{s}} \quad & f(\mathbf{s}) \\ \text{s.t.} \quad & \boldsymbol{\theta} \in \mathcal{K}(\hat{\mathbf{x}}, \tilde{\mathcal{X}}) \\ & \boldsymbol{\theta} \in \boldsymbol{\Theta}(\mathbf{s}) \\ & \mathbf{s} \in \mathcal{S}, \end{aligned} \tag{3}$$

and a subproblem that computes new solutions $\mathbf{x} \in \mathcal{X}$ which are then added to a growing set $\tilde{\mathcal{X}} \subset \mathcal{X}$ (Bodur et al. 2022). These algorithms will eventually converge at an optimal solution because of the observation that $\tilde{\mathcal{X}} = \text{ext}(\mathcal{X})$ is an exact surrogate for \mathcal{X} in $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X})$, where $\text{ext}(\mathcal{X})$ represents the set of extreme points of the convex hull of \mathcal{X} ; in the worst case, the algorithm can terminate after cycling through all points in $\text{ext}(\mathcal{X})$. However, because $\mathcal{K}(\hat{\mathbf{x}}, \tilde{\mathcal{X}})$ defines an outer approximation of $\mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$, standard CP approaches suffers from a fundamental drawback in that objective vectors generated by (3) are not feasible to (IO) until the algorithm terminates at the optimal solution, a process which can take an arbitrarily long time; see Sections 1 and 2 for additional discussion.

3.2.2. Approximation II: the linear relaxation of \mathcal{X} . Our paper is motivated by how we can best use the linear relaxation of \mathcal{X} , denoted as \mathcal{X}^{LR} . Note that since $\mathcal{X} \subseteq \mathcal{X}^{\text{LR}}$, $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}})$ contains a superset of the hyperplanes that define $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X})$, which implies that $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}) \subseteq \mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$ is inner approximation of the inverse-feasible region. Furthermore, since \mathcal{X}^{LR} is a convex polyhedral set, we can derive a simple polyhedral reformulation of the cone $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}})$, as shown below.

Observation 2. Given $\mathcal{X}^{\text{LR}} := \{\mathbf{x} \in \mathbb{R}^n \mid \mathbf{A}\mathbf{x} \leq \mathbf{b}\}$, we can reformulate $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}})$ as

$$\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}) = \{\boldsymbol{\theta} \mid \mathbf{p}^\top (\mathbf{A}\hat{\mathbf{x}} - \mathbf{b}) = 0, \mathbf{p}^\top \mathbf{A} = \boldsymbol{\theta}, \mathbf{p} \geq \mathbf{0}\}. \tag{4}$$

Note that if (FO) is a linear program, meaning that $\mathcal{X} = \mathcal{X}^{\text{LR}}$, then $\mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$ is equivalent to (4).

Equation (4) is derived using standard complementary slackness conditions. Note that for any index i , $p_i = 0$ if $\mathbf{a}_i^\top \hat{\mathbf{x}} - b_i \neq 0$, and $p_i \geq 0$ if $\mathbf{a}_i^\top \hat{\mathbf{x}} - b_i = 0$. This implies that the set $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}})$ is defined as the conic hull of normal vectors \mathbf{a}_i that correspond to binding constraints at $\hat{\mathbf{x}}$, i.e.,

$$\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}) = \text{cone}(\{\mathbf{a}_i\}_{i \in \mathcal{I}}), \tag{5}$$

where \mathcal{I} denotes the indices of all constraints $\mathbf{a}_i^\top \hat{\mathbf{x}} = b_i$.

Observation 2 implies that $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}})$ is tractable and provides a straightforward inner approximation of $\mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$. Using $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}})$, we can then consider the following approximation of (IO),

$$\begin{aligned} \min_{\boldsymbol{\theta}, \mathbf{s}} \quad & f(\mathbf{s}) \\ \text{s.t.} \quad & \boldsymbol{\theta} \in \mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}) \\ & \boldsymbol{\theta} \in \boldsymbol{\Theta}(\mathbf{s}) \\ & \mathbf{s} \in \mathcal{S}, \end{aligned} \tag{6}$$

which if feasible, would generate a feasible solution to (IO).

While this approximation of (IO) is simple and tractable, its main limitation is that the solution that is generated by (6) may not be very good. As an example, consider Figure 1 which shows an extreme case where $\hat{\mathbf{x}}$ lies strictly in the interior of \mathcal{X}^{LR} . In this case, $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}) = \{\mathbf{0}\}$. More generally, the observation implies that the quality of the IFR approximation depends entirely on the tightness of the linear relaxation at $\hat{\mathbf{x}}$.

Refining the linear relaxation \mathcal{X}^{LR} can improve the approximation of $\mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$ given by $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}})$. Equation (5) implies that such refinements are only beneficial when we add new valid inequalities that are both binding at $\hat{\mathbf{x}}$ and strengthen \mathcal{X}^{LR} (recall that a *valid inequality* is any inequality $\boldsymbol{\alpha}^\top \mathbf{x} \leq \beta$ that holds for all $\mathbf{x} \in \mathcal{X}$ (Cornuéjols 2008)). For example, Figure 2 illustrates how adding such valid inequalities to the linear relaxation in Figure 1 improves the corresponding IFR approximations, which are shown in the top right corner of each subfigure.

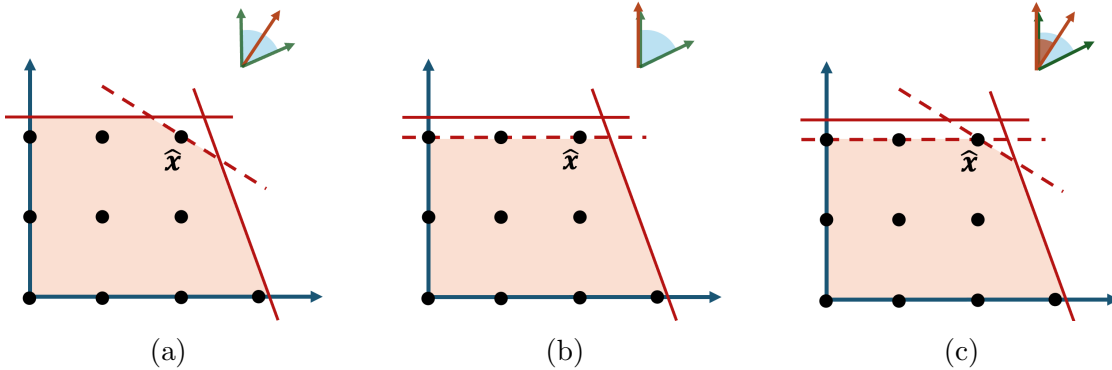


Figure 2 Tighter linear relaxations and their corresponding IFR approximations. In Figure 2a and 2b, the linear relaxation shown in Figure 1 is improved by adding a valid inequality, and each IFR approximation consists of a single (long orange) vector. Figure 2c illustrates when both inequalities are combined, leading to an IFR approximation defined by the cone between the two (long orange) vectors.

Nonetheless, the task of repeatedly generating these specific valid inequalities for discrete optimization problems is far from straightforward (e.g., Chvátal et al. 2013). Any procedure for generating new valid inequalities $\boldsymbol{\alpha}^\top \mathbf{x} \leq \beta$ that lead to improvements in $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}})$ would be of

APPROXIMATIONS OF \mathcal{X}	APPROXIMATIONS OF $\mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$	USE IN SOLVING IO PROBLEMS
$\tilde{\mathcal{X}} \subseteq \mathcal{X}$	outer approximation	CP Algorithms [iteratively increases size of $\tilde{\mathcal{X}}$] (Wang 2009, Bodur et al. 2022)
$\mathcal{X}^{\text{LR}} \supseteq \mathcal{X}$	inner approximation	N/A (see Section 3.2.2)
$\mathcal{X}^{\text{LR}}(\mathcal{W})$ (RLRs)	Pure-RLR: inner approx.	RLR-CP Algorithms iteratively reduces size of $\mathcal{X}^{\text{LR}}(\mathcal{W})$ while increasing size of $\tilde{\mathcal{X}} \subseteq \mathcal{X}$
	Mixed-RLR: “mixed” approx.	

Table 1 An overview of how different surrogates for \mathcal{X} can lead to different approximations of $\mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$, as well as how these surrogates can be used to solve inverse optimization problems.

the same difficulty as directly identifying inverse-feasible vectors $\alpha \in \mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$ lying outside of $\text{cone}(\{\mathbf{a}_i\}_{i \in \mathcal{I}})$. In Section 4, we circumvent this challenge by proposing a simple but versatile technique that leverages existing valid inequalities to generate better approximations of the IFR.

3.3. Summary and roadmap

The main challenge in solving (IO) lies in the difficulty of handling the inverse-feasible region. In Section 3.2, we introduced two simple approximation strategies for $\mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$ – or equivalently for $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X})$ – by replacing \mathcal{X} with a surrogate set $\tilde{\mathcal{X}}$. When this surrogate is defined as a subset of forward-feasible solutions ($\tilde{\mathcal{X}} \subseteq \mathcal{X}$), problem (IO) can be solved iteratively using a cutting plane algorithm, but at the cost of long solution times and a solution process that is entirely uninformative until it terminates. On the other hand, defining the surrogate as the linear relaxation \mathcal{X}^{LR} yields an approximation of (IO) that is much easier to solve, but may produce low-quality solutions with no systematic means of improving these solutions.

In the next section, we introduce the idea of *restricted linear relaxations* (RLRs), a simple yet powerful concept that modifies the linear relaxation to generate richer approximations of the inverse-feasible region (Section 4). We then use RLRs to develop a generalized framework for solving IO problems, where RLRs serve as a bridge between cutting plane algorithms and linear relaxations (Section 5). This framework overcomes the inherent limitations of each approach while retaining their respective strengths. For clarity, Table 1 summarizes the main ideas discussed thus far and serves as a reference for the material that follows.

Remark 3. We assume that $\hat{\mathbf{x}}$ lies on the boundary of \mathcal{X}^{LR} , or equivalently, that given any problem (FO), we know of at least one constraint or valid inequality that is binding at $\hat{\mathbf{x}}$. This assumption is

justified for two reasons. First, in most real-world problems, $\hat{\mathbf{x}}$ will naturally bind at a large number of constraints. Equality constraints present in (FO) are a simple example, as are variable bounds, e.g., if \mathcal{X} contains even a single binary variable ($x_i \in \{0, 1\}$), then any $\mathbf{x} \in \mathcal{X}$ is on the boundary of \mathcal{X}^{LR} , and the corresponding constraints defining this boundary (e.g., $x_i \geq 0$ or $x_i \leq 1$) would be sufficient. Second, in sensitivity and post-optimality analysis problems, $\hat{\mathbf{x}}$ is the optimal solution to a nominal objective vector $\hat{\boldsymbol{\theta}}$. By definition, $\hat{\boldsymbol{\theta}}^\top \mathbf{x} \leq \hat{\boldsymbol{\theta}}^\top \hat{\mathbf{x}}$ would define such a valid inequality.

4. IFR Approximations with Restricted Linear Relaxations

In this section, we introduce the concept restricted linear relaxations (RLRs) as a new technique for generating approximations of the inverse-feasible region. Loosely speaking, RLRs are defined by taking certain valid inequalities and modifying their right-hand-side values so that they are no longer valid. We show that these seemingly simple modifications yield a versatile approximation framework with desirable properties. RLRs will then be leveraged again in Section 5 to develop a richer family of solution methods for inverse optimization.

4.1. Definitions

We first present a general definition for restricted linear relaxations, before further classifying them into two types. An illustrative example is provided afterwards.

4.1.1. General definition of an RLR. We begin by introducing the term *restricting inequality*, which is defined with respect to $\hat{\mathbf{x}} \in \mathcal{X}$. Specifically, given an inequality $\mathbf{w}^\top \mathbf{x} \leq \tau$, we say it is a restricting inequality if $\mathbf{w}^\top \hat{\mathbf{x}} > \tau$ and $\mathbf{w} \in \mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$.

Definition 1. Given a linear relaxation \mathcal{X}^{LR} , a restricted linear relaxation $\mathcal{X}^{\text{LR}}(\mathcal{W})$ is defined as

$$\mathcal{X}^{\text{LR}}(\mathcal{W}) := \{\mathbf{x} \in \mathcal{X}^{\text{LR}} \mid \mathbf{x} \in \mathcal{W}\} \quad (7)$$

where $\mathcal{W} := \{\mathbf{x} \in \mathbb{R}^n \mid \mathbf{W}\mathbf{x} \leq \boldsymbol{\tau}\}$ is composed of only restricting inequalities for $(\hat{\mathbf{x}}, \mathcal{X})$.

By definition, the set of restricting inequalities in \mathcal{W} are not valid for the linear relaxation \mathcal{X}^{LR} , as they exclude the feasible solution $\hat{\mathbf{x}} \in \mathcal{X}$. Nonetheless, they are closely related to valid inequalities for \mathcal{X}^{LR} . Specifically, given any valid inequality $\mathbf{w}^\top \mathbf{x} \leq \bar{\tau}$ that is tight at $\hat{\mathbf{x}}$, meaning that $\bar{\tau} = \mathbf{w}^\top \hat{\mathbf{x}}$, we can define a restricting inequality $\mathbf{w}^\top \mathbf{x} \leq \tau$ simply by setting $\tau < \bar{\tau}$. For this reason, we refer to $\eta := \bar{\tau} - \tau \geq 0$ as the *depth* of the restricting inequality $\mathbf{w}^\top \mathbf{x} \leq \tau$.

4.1.2. Pure and mixed-RLRs. To better describe their mathematical properties, we classify RLRs into two types by distinguishing between whether the restricting inequalities satisfy an additional condition. Specifically, given a $\mathbf{w}^\top \mathbf{x} \leq \tau$, we say it is a *separating* restricting inequality if it satisfies the additional condition that

$$\mathbf{w}^\top \mathbf{z} \leq \tau \quad \forall \mathbf{z} \in \text{ext}(\mathcal{X}) \setminus \{\hat{\mathbf{x}}\}, \quad (8)$$

while any other restricting inequality is considered *non-separating*. Note that the term comes from the fact that when (8) is satisfied, $\mathbf{w}^\top \mathbf{x} = \tau$ can be considered as a separating hyperplane between $\hat{\mathbf{x}}$ and all (other) extreme points $\text{ext}(\mathcal{X})$. For brevity, we refer to these inequalities simply as separating and non-separating inequalities.

Definition 2. Any restricted linear relaxation $\mathcal{X}^{\text{LR}}(\mathcal{W})$ can be classified into one of the following:

- *Pure-RLR*: \mathcal{W} contains only separating inequalities.
- *Mixed-RLR*: \mathcal{W} contains at least one non-separating inequality.

Example 2. Figure 3 illustrates three RLRs, generated simply by decreasing the right-hand-side value of the valid inequalities in Figures 2a and 2b. Specifically, Figure 3a illustrates a pure-RLR, where the valid inequality from Figure 2a is made into a separating inequality. Figure 3b illustrates a mixed-RLR, where the depth of the same inequality is increased until it becomes a non-separating inequality that excludes the extreme point $\mathbf{x}^e \in \text{ext}(\mathcal{X})$ shown. Finally, Figure 3c illustrates another mixed-RLR defined by shifting the valid inequality from Figure 2b, which also excludes \mathbf{x}^e .

4.2. Foundations and main insights

In this section, we lay the foundations of RLRs by establishing their basic properties and by illustrating their role in approximating the inverse-feasible region. Specifically, given a restricted linear relaxation $\mathcal{X}^{\text{LR}}(\mathcal{W})$, we consider the approximation of $\mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$ given by $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}))$. Note that since $\mathcal{X}^{\text{LR}}(\mathcal{W}) \subseteq \mathcal{X}^{\text{LR}}$, it follows from (2) that $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}) \subseteq \mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}))$. Furthermore, since $\mathcal{X}^{\text{LR}}(\mathcal{W})$ is a polyhedral set, we can also derive a tractable polyhedral description for $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}))$, much like Observation 2.

Proposition 1. Let $\mathcal{X}^{\text{LR}}(\mathcal{W}) := \{\mathbf{x} \in \mathbb{R}^n \mid \mathbf{A}\mathbf{x} \leq \mathbf{b}, \mathbf{W}\mathbf{x} \leq \boldsymbol{\tau}\}$ denote a restricted linear relaxation. Then, $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}))$ can be reformulated as

$$\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W})) = \{\boldsymbol{\theta} \mid \boldsymbol{\theta}^\top \hat{\mathbf{x}} \geq \mathbf{p}^\top \mathbf{b} + \mathbf{z}^\top \boldsymbol{\tau}, \mathbf{A}^\top \mathbf{p} + \mathbf{W}^\top \mathbf{z} = \boldsymbol{\theta}, \mathbf{p} \geq \mathbf{0}, \mathbf{z} \geq \mathbf{0}\}. \quad (9)$$

In Figure 3, the approximation $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}))$ generated by each RLR is shown in the top right corner of each subfigure, overlaid onto the inverse-feasible region $\mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$. First, we note that the approximations in each subfigure captures a sizeable portion of $\mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$, in contrast to those of Figures 2a and 2b, which contains only a single inverse-feasible vector. Second, we observe that in Figure 3a, $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}^a))$ forms a valid inner approximation of $\mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$, whereas $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}^b))$ and $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}^c))$ in Figures 3b and 3c include a small region of vectors that are not inverse-feasible. We now present two results that explain these two observations.

Theorem 1. For any restricted linear relaxation $\mathcal{X}^{\text{LR}}(\mathcal{W})$, the following three conditions hold:

- $\dim(\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}))) = n$,

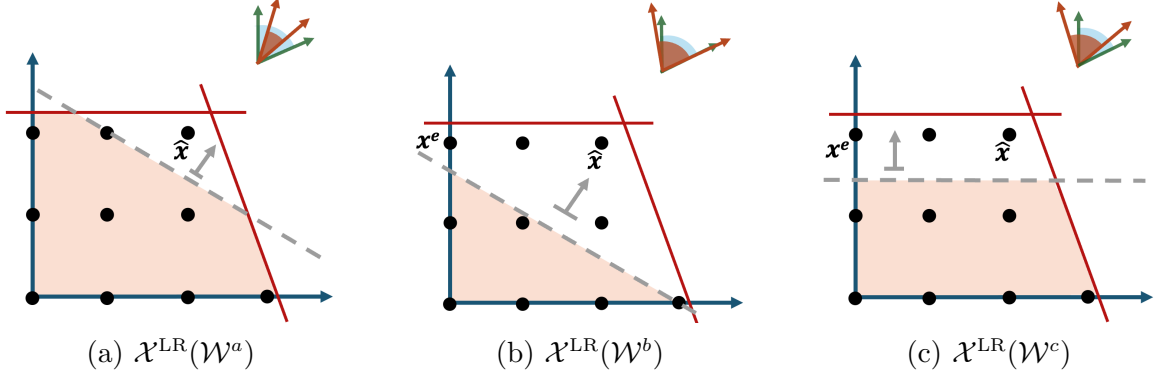


Figure 3 Three restricted linear relaxations, and their corresponding IFR approximations overlaid onto the true IFR. Figure (a) illustrates a pure-RLR generated with a separating inequality, while (b) and (c) illustrate mixed-RLRs generated with a non-separating inequality.

- $w \in \text{int}(\mathcal{K}(\hat{x}, \mathcal{X}^{\text{LR}}(\mathcal{W})))$ for all $w \in \mathbf{W}$,
- $\dim(\mathcal{K}(\hat{x}, \mathcal{X}^{\text{LR}}(\mathcal{W})) \cap \mathcal{R}_{\text{inv}}(\hat{x}, \mathcal{X})) = \dim(\mathcal{R}_{\text{inv}}(\hat{x}, \mathcal{X}))$.

Theorem 1 highlights that $\mathcal{K}(\hat{x}, \mathcal{X}^{\text{LR}}(\mathcal{W}))$ is a full dimensional cone in \mathbb{R}^n encompassing the vectors w_i in its interior. Furthermore, its intersection with $\mathcal{R}_{\text{inv}}(\hat{x}, \mathcal{X})$ spans the same subspace as $\mathcal{R}_{\text{inv}}(\hat{x}, \mathcal{X})$, implying that RLRs always yield a substantive approximation of the inverse-feasible region. Figures 3a, 3b and 3c show that each approximation captures a sizeable portion of $\mathcal{R}_{\text{inv}}(\hat{x}, \mathcal{X})$, in contrast to those generated by the linear relaxations in Figures 2a and 2b.

Theorem 2. For any restricted linear relaxation $\mathcal{X}^{\text{LR}}(\mathcal{W})$, the condition

$$\mathcal{K}(\hat{x}, \mathcal{X}^{\text{LR}}(\mathcal{W})) \subseteq \mathcal{R}_{\text{inv}}(\hat{x}, \mathcal{X}) \quad (10)$$

is satisfied if and only if $\mathcal{X}^{\text{LR}}(\mathcal{W})$ is a pure-RLR.

Theorem 2 specifies that $\mathcal{K}(\hat{x}, \mathcal{X}^{\text{LR}}(\mathcal{W}))$ is an inner approximation of $\mathcal{R}_{\text{inv}}(\hat{x}, \mathcal{X})$ if and only if \mathcal{W} consists of only separating inequalities. This explains why the approximation in Figure 3a lies strictly within $\mathcal{R}_{\text{inv}}(\hat{x}, \mathcal{X})$, whereas those of Figure 3b and 3c contain inverse-infeasible vectors.

Practically, RLRs provide a mechanism to construct IFR approximations of varying sizes, where smaller RLRs always yield approximations that are at least as large – and in most cases strictly larger – than those generated by larger RLRs. This follows directly from the definition of $\mathcal{K}(\hat{x}, \cdot)$, where, for any $\mathcal{X}' \subset \mathcal{X}'' \subset \mathcal{X}^{\text{LR}}$, it must be true that $\mathcal{K}(\hat{x}, \mathcal{X}^{\text{LR}}) \subseteq \mathcal{K}(\hat{x}, \mathcal{X}'') \subseteq \mathcal{K}(\hat{x}, \mathcal{X}')$. For example, when the depth of the restricting inequality is increased between Figures 3a and 3b, the size of the corresponding approximation $\mathcal{K}(\hat{x}, \mathcal{X}^{\text{LR}}(\mathcal{W}))$ becomes larger, encompassing more of $\mathcal{R}_{\text{inv}}(\hat{x}, \mathcal{X})$ but also including vectors lying outside of $\mathcal{R}_{\text{inv}}(\hat{x}, \mathcal{X})$.

Up to this point, we have relied on simple illustrative examples to develop intuition for the use of RLRs. We now present a richer example to better highlight the nuances and trade-offs associated

with changes in the depth of a restricting inequality and, more generally, the size of an RLR. This example, shown below, also provides a preview of the post-optimality analysis problems that will be the focus of Section 6.

Example 3. Consider the classic capital budgeting problem, which aims to determine an optimal set of project investments that maximizes the total revenue from the selected projects (Petersen 1967). This problem is formulated as

$$\begin{aligned} \max_{\mathbf{x}} \quad & \hat{\boldsymbol{\theta}}^\top \mathbf{x} \\ \text{s.t.} \quad & \mathbf{x} \in \mathcal{X} := \{\mathbf{x} \in \{0,1\}^n \mid \mathbf{A}\mathbf{x} \leq \mathbf{b}\}, \end{aligned} \quad (11)$$

where $\hat{\boldsymbol{\theta}}$ denotes the predicted revenue of the projects, $\mathbf{A}\mathbf{x} \leq \mathbf{b}$ represents the budget constraints, and $\mathbf{x} \in \mathcal{X}$ indicates the binary investment decisions. Let $\hat{\mathbf{x}} \in \mathcal{X}$ denote an optimal solution to this problem. For this solution, recall that we can generate exact sensitivity ranges $[\hat{\theta}_i + \Delta_i^-, \hat{\theta}_i + \Delta_i^+]$ for each coefficient $\hat{\theta}_i$, holding all others fixed, by solving a pair of inverse optimization problems (models (1a) and (1b)). Using these exact ranges as a benchmark, we examine the ranges that are generated when we replace $\mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$ with approximations $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}))$ generated by different RLRs. We note that for binary programs, the sensitivity range for each coefficient is always unbounded in one direction, i.e., $\Delta_i^- = -\infty$ when $\hat{x}_i = 0$ and $\Delta_i^+ = \infty$ when $\hat{x}_i = 1$, and we thus consider only the range in the bounded direction (e.g., $[\hat{\theta}_i, \hat{\theta}_i + \Delta_i^+]$ when $\hat{x}_i = 0$).

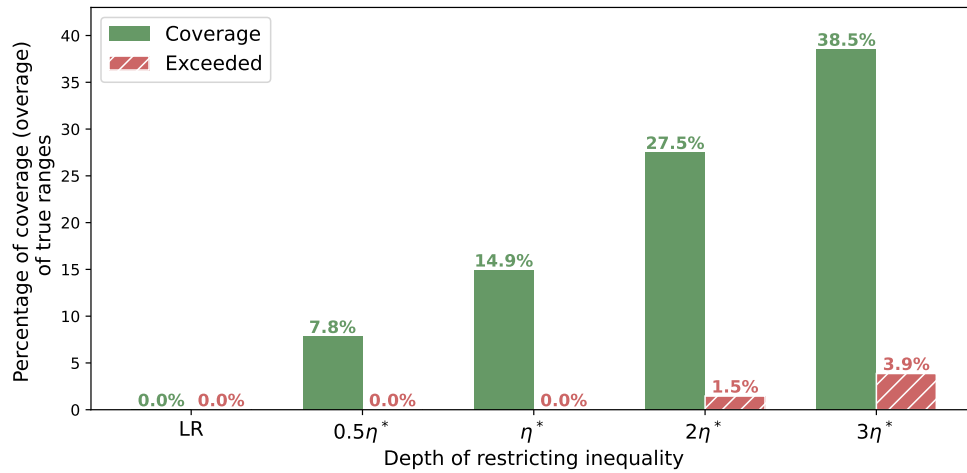


Figure 4 The percentage of total sensitivity ranges that are covered, and the amount that exceeds the range.

Figure 4 provides a summary of the sensitivity ranges generated for a well-known capital budgeting problem from Petersen (1967), containing 50 variables and 5 budget constraints. Specifically, we consider ranges generated by \mathcal{X}^{LR} , denoting the linear relaxation of (11), and by $\mathcal{X}^{\text{LR}}(\mathcal{W}) := \{\mathbf{x} \in \mathcal{X}^{\text{LR}} \mid \hat{\boldsymbol{\theta}}^\top \mathbf{x} \leq \hat{\boldsymbol{\theta}}^\top \hat{\mathbf{x}} - \eta\}$ for four different values of the depth parameter $\eta > 0$. In the figure,

η^* denotes the maximal depth for which $\hat{\theta}^\top \mathbf{x} \leq \hat{\theta}^\top \hat{\mathbf{x}} - \eta^*$ remains a separating inequality; methods to compute η^* are deferred to Section 4.3. Each green bar shows the percentage of the true sensitivity range captured by the approximation (referred to as “coverage”), aggregated over all 50 objective coefficients. Similarly, each red striped bar shows the percentage by which the approximations exceeded the true sensitivity ranges (referred to as “overapproximation”). The complete details of this numerical example can be found in Section A2.2 of the Electronic Companion.

There are several insights that we can derive from Example 3. First, the linear relaxation \mathcal{X}^{LR} fails to provide any informative approximation of the sensitivity ranges; in fact, as we illustrate in Section A2.2, this remains true even for a tighter linear relaxation. This highlights the inherent limitations of standard linear relaxations. Second, the restricting inequality with a depth of $0.5\eta^*$ and η^* yield sensitivity ranges that remain entirely within the true ranges, with a coverage of 7.8% and 14.9%, respectively. This demonstrates both the idea that separating inequalities generate inner approximation of the IFR (Theorem 2), and the value of computing deeper separating inequalities when possible. Third, the results for depths of $2\eta^*$ and $3\eta^*$ illustrate the trade-off between the level of coverage and overapproximation: increasing the depth results in approximations that encompass more of the inverse-feasible region while admitting a small region of inverse-infeasible vectors. The same trend is observed when moving from $2\eta^*$ and $3\eta^*$.

The examples above clearly illustrate why RLRs outperform standard linear relaxations. More importantly, they highlight the versatility of RLRs and the practical trade-offs that arise when considering RLRs of different sizes, which will be a central theme in the remainder of this paper:

- RLRs with shallower restricting inequalities are more likely to yield inner approximations of the inverse-feasible region, though such approximations may be overly conservative.
- RLRs with deeper restricting inequalities can capture larger portions of the inverse-feasible region, but may also include regions that are not inverse-feasible.

The next two subsections establish a more complete theoretical foundation for these observations. Specifically, Section 4.3 examines the computation of the maximal depth parameter η^* beyond which a valid or separating inequality becomes non-separating, while Section 4.4 formalizes conditions under which RLRs yield strictly better approximations of the IFR. The trade-off highlighted above is also central to understanding the characteristics of the solution framework presented in Section 5, and will be further explored in the numerical experiments of Section 6.

4.3. Separating and non-separating inequalities

In this section, we more carefully consider the computation of separating and non-separating inequalities, or more specifically, the depth threshold at which a valid or separating inequality

$\mathbf{w}^\top \mathbf{x} \leq \tau$ becomes a non-separating one. For a given vector $\mathbf{w} \in \mathbb{R}^n$, this task can, in principle, be reduced to solving the modified forward optimization problem

$$\tau^*(\mathbf{w}) = \max_{\mathbf{x}} \{ \mathbf{w}^\top \mathbf{x} \mid \mathbf{x} \in \text{ext}(\mathcal{X}) \setminus \{\hat{\mathbf{x}}\} \}. \quad (12)$$

By definition, $\mathbf{w}^\top \mathbf{x} \leq \tau$ for any $\tau < \tau^*(\mathbf{w})$ would define a non-separating inequality. Similarly, if $\tau^*(\mathbf{w}) < \mathbf{w}^\top \hat{\mathbf{x}}$, then $\mathbf{w}^\top \mathbf{x} \leq \tau$ for any $\tau \in [\tau^*(\mathbf{w}), \mathbf{w}^\top \hat{\mathbf{x}})$ would define a separating inequality; in this case, we refer to $\mathbf{w}^\top \mathbf{x} \leq \tau^*(\mathbf{w})$ as a *tight* separating inequality. Finally, if $\tau^*(\mathbf{w}) = \mathbf{w}^\top \hat{\mathbf{x}}$, this implies that no separating inequality can be generated.

The main challenge in solving Problem (12) is that the set of extreme points $\text{ext}(\mathcal{X})$ is not known a priori. If it were, $\mathcal{K}(\hat{\mathbf{x}}, \text{ext}(\mathcal{X}))$ would provide an exact description of the IFR, and there would be no need to solve this problem or generate approximations in the first place. As a proxy, we consider a formulation for generating separating inequalities that exclude only $\hat{\mathbf{x}}$ from \mathcal{X} , i.e.,

$$\max_{\mathbf{x}} \{ \mathbf{w}^\top \mathbf{x} \mid \mathbf{x} \in \mathcal{X} \setminus \{\hat{\mathbf{x}}\} \}. \quad (13)$$

We note that when $\mathcal{X} \subseteq \{0, 1\}^n$ is defined only by binary variables, Problems (12) and (13) are in fact equivalent, since $\mathcal{X} = \text{ext}(\mathcal{X})$. In this setting, we derive a simple mathematical program for solving them, shown in the next proposition.

Proposition 2. *When $\mathcal{X} \subseteq \{0, 1\}^n$, Problem (12) and (13) can be formulated and solved as*

$$\tau^*(\mathbf{w}) = \max_{\mathbf{x}} \quad \mathbf{w}^\top \mathbf{x} \quad (14a)$$

$$\text{s.t.} \quad \mathbf{x} \in \mathcal{X} \quad (14b)$$

$$\sum_{i \in J_1(\hat{x}_i)} x_i + \sum_{i \in J_0(\hat{x}_i)} (1 - x_i) \leq n - 1, \quad (14c)$$

where $J_1(\hat{x}_i)$ and $J_0(\hat{x}_i)$ denote the indices $i \in \{1, \dots, n\}$ for which \hat{x}_i is equal to 1 and 0, respectively.

When $\mathcal{X} \subseteq \mathbb{Z}^n$ contains general integer variables, we can also derive a similar mathematical program to solve Problem (13), albeit one that is larger in size and requires introducing additional auxiliary variables. For brevity, we provide this formulation in **Proposition A1** in the Electronic Companion. We note that in this more general integer case where $\mathcal{X} \neq \text{ext}(\mathcal{X})$, Problem (13) is not equivalent to Problem (12), which implies that Problem (13) can still be used to generate separating inequalities but they may not be tight.

Finally, we note that separating inequalities – and, by extension, pure-RLRs – can be defined only in certain contexts; specifically, they exist only when $\hat{\mathbf{x}} \in \text{ext}(\mathcal{X})$. This follows directly from the observation that if $\hat{\mathbf{x}} \notin \text{ext}(\mathcal{X})$, then it would not be possible to separate $\hat{\mathbf{x}}$ from $\text{ext}(\mathcal{X})$.

4.4. Conditions for improved approximations

In this final subsection, we derive mathematical conditions that further clarify when an RLR leads to a better approximation of the inverse-feasible region, relative to either \mathcal{X}^{LR} or another RLR. Our analysis focuses on pure-RLRs. This is both to simplify the exposition and because mixed-RLRs can largely be interpreted as extensions of pure-RLRs where the depth of separating inequalities are further increased until they become non-separating; hence, the insights established for pure-RLRs remain relevant for mixed-RLRs when interpreted through this lens. Finally, all results are presented in terms of the expansion of $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}))$. Since pure-RLRs lead to inner approximations of the inverse-feasible region (Theorem 2), any expansion of $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}))$ is equivalent to capturing a larger portion of this region.

4.4.1. Improvements on \mathcal{X}^{LR} . We first present conditions under which the approximation $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}))$ generated by an RLR is better than the approximation $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}})$ generated by \mathcal{X}^{LR} . Without loss of generality, we focus on RLRs defined by a single separating inequality.

Theorem 3. *Let $\mathcal{X}^{\text{LR}}(\mathcal{W}) := \{\mathbf{x} \in \mathcal{X}^{\text{LR}} \mid \mathbf{w}^\top \mathbf{x} \leq \tau\}$ denote an RLR defined using a separating inequality $\mathbf{w}^\top \mathbf{x} \leq \tau$. Then,*

$$\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}) \subset \mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W})) \quad (15)$$

if and only if there exists a $\mathbf{x} \in \text{ext}(\mathcal{X}^{\text{LR}}) \setminus \{\hat{\mathbf{x}}\}$ where $\mathbf{w}^\top \mathbf{x} > \tau$.

Put simply, Theorem 3 states that if the inequality $\mathbf{w}^\top \mathbf{x} \leq \tau$ is able to cut off any extreme point of \mathcal{X}^{LR} excluding $\hat{\mathbf{x}}$, then $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}))$ will yield a strictly larger inner approximation of $\mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$ than $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}})$. The intuition is that extreme points of \mathcal{X}^{LR} define the corresponding approximation, i.e., $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}) = \mathcal{K}(\hat{\mathbf{x}}, \text{ext}(\mathcal{X}^{\text{LR}}))$, and eliminating at least one with a separating inequality will enlarge the approximation but preserve its validity as an inner approximation.

Proposition 3. *Let \mathbf{w} be a vector on the boundary of $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}})$. If \mathbf{w} can generate a separating inequality $\mathbf{w}^\top \mathbf{x} \leq \tau$, where $\mathcal{X}^{\text{LR}}(\mathcal{W}) := \{\mathbf{x} \in \mathcal{X}^{\text{LR}} \mid \mathbf{w}^\top \mathbf{x} \leq \tau\}$, then*

$$\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}) \subset \mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W})).$$

Proposition 3 provides an explanation for one of the observations made in Example 3. In this example, the nominal objective vector $\hat{\boldsymbol{\theta}}$ in fact defines a facet of (a tighter version of) \mathcal{X}^{LR} and thus lies on the boundary of $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}})$ (based on Observation 2). In the example, we saw that $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}})$ did not provide any informative sensitivity ranges, but the resulting $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}))$ did. More broadly, this observation illustrates that the facets of \mathcal{X}^{LR} that are tangent to $\hat{\mathbf{x}}$ – the normal vectors of which lie on the boundary of $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}})$ – serve as good candidates for constructing restricting inequalities. Additional discussion on this example can be found in Section A2.2.

4.4.2. Improvements across RLRs. We now derive conditions for when an IFR approximation improves between two RLRs. We first focus on the depth of a single separating inequality. We show that if the addition of a separating inequality to \mathcal{X}^{LR} improves the corresponding IFR approximation, then increasing it's depth will result in an even better approximation.

Theorem 4. *Let $\mathcal{X}^{\text{LR}}(\mathcal{W}_1) := \{\mathbf{x} \in \mathcal{X}^{\text{LR}} \mid \mathbf{w}^\top \mathbf{x} \leq \tau_1\}$ denote an RLR where $\mathbf{w}^\top \mathbf{x} \leq \tau_1$ is a separating inequality that is not tight (see Section 4.3). If $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}) \subset \mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}_1))$, then, for any $\mathcal{X}^{\text{LR}}(\mathcal{W}_2) := \{\mathbf{x} \in \mathcal{X}^{\text{LR}} \mid \mathbf{w}^\top \mathbf{x} \leq \tau_2\}$ where $\tau_2 < \tau_1$,*

$$\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}_1)) \subset \mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}_2)).$$

While the discussion up to this point has focused on RLRs defined by a single separating inequality, we can extend the analysis to the case of multiple separating inequalities. Specifically, for a given $\mathcal{X}^{\text{LR}}(\mathcal{W}^1)$, consider a refined $\mathcal{X}^{\text{LR}}(\mathcal{W}^2) := \{\mathbf{x} \in \mathcal{X}^{\text{LR}}(\mathcal{W}^1) \mid \mathbf{w}^\top \mathbf{x} \leq \tau\}$ that has an additional separating inequality. Theorem 3 naturally extends to this setting with only minor modifications.

Theorem 5. *Let $\mathcal{X}^{\text{LR}}(\mathcal{W}^2) := \{\mathbf{x} \in \mathcal{X}^{\text{LR}}(\mathcal{W}^1) \mid \tilde{\mathbf{w}}^\top \mathbf{x} \leq \tilde{\tau}\}$ denote an RLR defined by adding a separating inequality $\tilde{\mathbf{w}}^\top \mathbf{x} \leq \tilde{\tau}$ to an existing $\mathcal{X}^{\text{LR}}(\mathcal{W}^1)$. Then,*

$$\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}^1)) \subset \mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}^2)) \tag{16}$$

if and only if there exists a $\mathbf{x} \in \text{ext}(\text{conv}(\mathcal{X}^{\text{LR}}(\mathcal{W}^1) \cup \{\hat{\mathbf{x}}\}))$ where $\tilde{\mathbf{w}}^\top \mathbf{x} > \tilde{\tau}$.

The basic intuition behind this result is that $\text{conv}(\mathcal{X}^{\text{LR}}(\mathcal{W}^1) \cup \{\hat{\mathbf{x}}\})$ is itself a valid linear relaxation for \mathcal{X} , in which case Theorem 3 can be applied directly to it.

4.4.3. Summary and implications. The results above formalize several conditions under which RLRs yield improved approximations of the inverse-feasible region. Taken together, they convey a simple but important message: larger portions of the inverse-feasible region can be captured either by increasing the depth of restricting inequalities or by adding more of them. This intuition serves as the conceptual basis for the IO solution methods developed in the next section.

5. RLR-Driven Algorithms for Inverse Optimization

In this section, we leverage RLRs to introduce a new paradigm for solving inverse optimization problems. Specifically, we propose a nested decomposition approach that integrates RLRs with a cutting plane subroutine, which yields a versatile framework that directly addresses the main limitations of standalone cutting plane algorithms. We first describe the approach in Section 5.1, then examine its properties and features in Section 5.2.

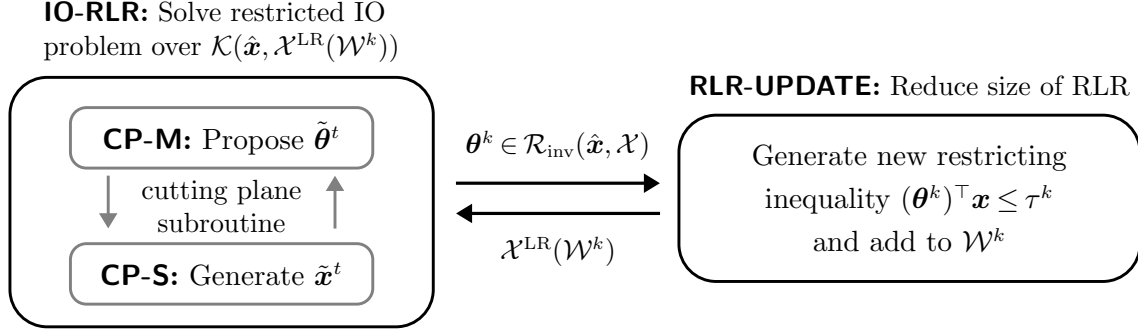


Figure 5 Visualization of the RLR-driven decomposition framework for inverse optimization.

5.1. Description of the algorithm

We begin with a high-level summary of the approach, which we refer to as the RLR-CP Algorithm. This algorithm iteratively leverages RLRs to restrict the solution space of the inverse optimization problem. Specifically, within each iteration, a cutting plane subroutine refines the space by eliminating regions that are not inverse-feasible, while across iterations, the space is progressively expanded by adding restricting inequalities to the linear relaxation \mathcal{X}^{LR} . The degree at which the solution space is expanded is determined by the depth of the restricting inequalities.

The complete details of the RLR-CP Algorithm is presented in Algorithm 1, and a schematic overview is provided in Figure 5. At its core, the RLR-CP Algorithm follows a nested decomposition structure. In each iteration of the outer loop, indexed by k , the algorithm first solves IO-RLR, which defines a restricted version of Problem (IO), i.e.,

$$\min_{\theta, s} f(s) \quad (17a)$$

$$\text{s.t. } \theta \in \mathcal{K}(\hat{x}, \mathcal{X}^{LR}(\mathcal{W}^k)) \quad (17b)$$

$$\theta \in \mathcal{R}_{inv}(\hat{x}, \mathcal{X}) \quad (17c)$$

$$\theta \in \Theta(s) \quad (17d)$$

$$s \in \mathcal{S}. \quad (17e)$$

Solving this problem generates an inverse-feasible vector $\theta^k \in \mathcal{K}(\hat{x}, \mathcal{X}^{LR}(\mathcal{W}^k))$. If $f(s^k)$ does not strictly improve on $f(s^{k-1})$, then the RLR-CP Algorithm terminates and returns (θ^k, s^k) as a solution to (IO). Otherwise, the subproblem RLR-UPDATE uses θ^k to generate a restricting inequality $(\theta^k)^\top x \leq \tau^k$, which is added to the existing set of restricting inequalities \mathcal{W}^k , creating \mathcal{W}^{k+1} . Recall that this restricting inequality can also be written as

$$(\theta^k)^\top x \leq (\theta^k)^\top \hat{x} - \eta^k \quad (18)$$

where $\eta^k > 0$ denotes the depth of this inequality; η^k will be the focus of Sections 5.2 and 5.3.

A cutting plane subroutine is used to solve (17). This subroutine iterates between its own master problem CP-M and subproblem CP-S. At each iteration of this inner loop, indexed by t , CP-M

Algorithm 1: RLR-CP Algorithm for Inverse Optimization

```

1 Initialization:  $k = 0, t = 0, \text{OptGap} = \infty, \Delta^0 = \infty, u^{-1} = \infty, \mathcal{W}^0 = \mathbb{R}^n, \tilde{\mathcal{X}} = \{\emptyset\}$ 
2 while  $\Delta^k > 0$  or  $\text{OptGap} \leq \varepsilon$  do
3   IO-RLR: repeat
4     CP-M: Solve
          
$$\begin{aligned} \min_{\theta, s} \quad & f(s) \\ \text{s.t.} \quad & \theta \in \mathcal{K}(\hat{x}, \mathcal{X}^{\text{LR}}(\mathcal{W}^k)) \\ & \theta^\top (\hat{x} - x) \geq 0 \quad \forall x \in \tilde{\mathcal{X}} \\ & \theta \in \Theta(s) \\ & s \in \mathcal{S}. \end{aligned}$$

5     Let  $(\tilde{\theta}^t, \tilde{s}^t)$  denote the optimal solution.
6     CP-S: Generate  $\tilde{x}^t \in \arg \max\{(\tilde{\theta}^t)^\top x \mid x \in \mathcal{X}\}$ . If  $(\tilde{\theta}^t)^\top (\tilde{x}^t - \hat{x}) = 0$ , set  $\tilde{x}^t \leftarrow \emptyset$ .
7      $\tilde{\mathcal{X}} \leftarrow \tilde{\mathcal{X}} \cup \{\tilde{x}^t\}; t \leftarrow t + 1$ 
8   until  $\tilde{x}^t = \emptyset$ ;
9   – Parameter updates:  $(\theta^k, s^k) \leftarrow (\tilde{\theta}^t, \tilde{s}^t); u^k \leftarrow f(s^k); \Delta^k \leftarrow u^{k-1} - u^k$ ;
10  – Compute dual bound  $\ell^k$  and optimality gap:
11    Solve  $\text{IO}(\hat{x}, \tilde{\mathcal{X}})$ , let  $\ell^k$  denote optimal objective value.
12     $\text{OptGap} \leftarrow \frac{u^k - \ell^k}{\ell^k}$ 
13  RLR-UPDATE: Generate depth parameter  $\eta^k > 0$  and set  $\tau^k \leftarrow (\theta^k)^\top \hat{x} - \eta^k$ 
14     $\mathcal{W}^{k+1} \leftarrow \{x \in \mathcal{W}^k \mid (\theta^k)^\top x \leq \tau^k\}$ ;
15     $k \leftarrow k + 1; t \leftarrow 0$ 
16 end
17 return  $(\theta^k, s^k)$ 

```

generates a candidate vector $\tilde{\theta}^t$ by solving a relaxed version of (17) where $\mathcal{R}_{\text{inv}}(\hat{x}, \mathcal{X})$ is replaced with $\mathcal{K}(\hat{x}, \tilde{\mathcal{X}})$ for some subset of enumerated forward-feasible solutions $\tilde{\mathcal{X}} \subset \mathcal{X}$. The subproblem CP-S then tests if this candidate vector $\tilde{\theta}^t$ is inverse-feasible. If it is not, this implies the existence of a solution $\tilde{x}^t \in \mathcal{X}$ that is strictly better than \hat{x} under $\tilde{\theta}^t$. This solution \tilde{x}^t is then added to $\tilde{\mathcal{X}}$ and serves as a new constraint for the next iteration of CP-M. We refer to Bodur et al. (2022) for a detailed overview of different subproblems CP-S that can be used.

5.2. Analysis and notable features

The RLR-CP Algorithm offers a versatile framework for inverse optimization. It can be viewed both as a generalization of the standard cutting plane approach, and as a method capable of operating independently from it. Specifically, as its name indicates, the algorithm integrates two

components, each of which, in the extreme case, could function without the other. At one extreme, if we were to initialize $\mathcal{W}^0 = \emptyset$ as an empty set, then $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}^0)) = \mathbb{R}^n$ would not restrict the solution space in any way, and the RLR-CP Algorithm would simply reduce to the standard cutting plane approach (i.e., the optimal solution would be generated in the first iteration of IO-RLR). On the other hand, if we initialize $\mathcal{W}^0 := \mathbb{R}^n$ without restricting inequalities, or if \mathcal{W}^k were to be defined by only separating inequalities, then the corresponding $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}^k))$ is always an inner approximation of $\mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$, rendering the cutting plane subroutine redundant for those iterations (i.e., CP-S would not generate any cuts).

Algorithmic flexibility. Given these two extremes, the depth parameter η^k provides the flexibility to operate between them, which in turn affects both the characteristics and the performance of the RLR-CP Algorithm. When the value of η^k is large, the algorithm behaves more like the standard cutting plane algorithm, since the corresponding set $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}^k))$ is relatively large. As a result, solving IO-RLR will typically require more time, but can yield significantly better solutions when solved. Conversely, when η^k is small, the algorithm will generally spend less time within the cutting plane subroutine, which results in more frequent, albeit smaller, improvements through iterations of the outer loop. In our numerical experiments, we will show that the RLR-CP Algorithm consistently achieves better performance when operating between these two extremes.

Primal and dual bounds. By design, the RLR-CP Algorithm provides a framework that can simultaneously improve primal (i.e., inverse-feasible) solutions and dual bounds for the IO problem. Specifically, the outer iterations drive improvements in the primal solutions, while the cutting plane subroutine generates feasible solutions in \mathcal{X} that can be used to strengthen the dual bound. The parameter η^k can again be viewed as a lever that balances the time spent on each. As shown in Algorithm 1, this feature of the RLR-CP Algorithm permits the definition of a pre-determined optimality gap at which the solution process can be terminated. This provides a significant advantage over the standalone cutting plane approach, which returns a primal solution only at termination. In our numerical experiments, we demonstrate that the RLR-CP Algorithm can generate near-optimal solutions in far less time than is required to compute the optimal solution.

Convergence to optimality. The previous two paragraphs highlight the effect of η^k on the performance of the RLR-CP Algorithm. We now state a simple, sufficient condition on the value of η^k that ensures that the algorithm will converge to an optimal solution of Problem (IO). As stated previously, we restrict our attention to IO problems where $f(\mathbf{s})$ and \mathcal{S} are assumed to be convex.

Proposition 4. *Suppose Problem (IO) is convex and bounded. If the master problem (17) is feasible in $k = 0$, and if $\eta^k > 0 \forall k \geq 0$, then Algorithm 1 returns an optimal solution to (IO).*

The conditions required for Proposition 4 are minimal. Put simply, the proposition states that once the first restricting inequality can be generated, the algorithm is guaranteed to return an optimal solution provided that every subsequent restricting inequality added has non-trivial depth.

5.3. Choosing the depth parameter

As discussed above, the characteristics and performance of the RLR-CP Algorithm are largely determined by the choice of η^k at each iteration. Below, we outline two general frameworks for specifying η^k . As a general guideline, when solutions to (IO) must be generated under strict time limits, smaller time limits should generally be paired with smaller depth parameters, which tend to produce faster – albeit more incremental – improvements in inverse-feasible solutions. Conversely, larger depths can provide greater improvements at the expense of increased computation time.

Static policy. We first consider a static policy, which refers to a pre-determined and fixed rule for choosing η^k . Such policies are simple and easy to implement, making it a natural starting point for understanding the behavior of the RLR-CP Algorithm; for this reason, they will be extensively tested and examined in Section 6. The most straightforward static policy is to define η^k as fixed fraction or multiple of $|(\boldsymbol{\theta}^k)^\top \hat{\mathbf{x}}|$. Specifically, consider

$$\eta^k = \tilde{\eta} \cdot |(\boldsymbol{\theta}^k)^\top \hat{\mathbf{x}}| \quad (19)$$

where $\tilde{\eta} > 0$ defines some fixed (i.e., static) value. This depth selection method is immune to the arbitrary scaling of $\boldsymbol{\theta}$ values, which is important since $\mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$ is a cone. Furthermore, in problem settings where $\mathcal{X} \subset \mathbb{R}_+^n$ and $\boldsymbol{\theta}^k \in \mathbb{R}_+^n$ is restricted to lie in the nonnegative orthant, $\tilde{\eta} \in [0, 1)$ is sufficient for operating between the two extremes discussed in Section 5.2.

We can also replace $|(\boldsymbol{\theta}^k)^\top \hat{\mathbf{x}}|$ in (19) with $|(\boldsymbol{\theta}^k)^\top \hat{\mathbf{x}} - \tau^*(\boldsymbol{\theta}^k)|$, where $\tau^*(\boldsymbol{\theta}^k)$ is the solution to Problem (13) defining the depth of a tight separating inequality. Example 3 in Section 4.1 illustrated the coverage of the inverse-feasible region when the depth parameter is defined based on a multiple of $|(\boldsymbol{\theta}^k)^\top \hat{\mathbf{x}} - \tau^*(\boldsymbol{\theta}^k)|$. Note that the static policy based on $\tau^*(\boldsymbol{\theta}^k)$ can only be implemented for pure binary or integer forward problems (for more details, see Section 4.3).

Adaptive policy. The framework defined by the RLR-CP Algorithm also allows us to introduce a policy that adjusts the depth parameter dynamically based on the algorithm's real-time performance. This both mitigates the algorithm's sensitivity to the initial depth values specified under the static policy, and enhances its performance through a more automated parameter selection process. Due to space considerations, we defer its presentation and discussion to Section A3.1 of the Electronic Companion. To the best of our knowledge, this approach constitutes the first adaptive solution method proposed for inverse optimization.

6. Applications and Computational Insights

In this section, we first outline how inverse optimization can be applied to the sensitivity and post-optimality analysis of modern optimization problems (Section 6.1). We then focus on one such application, where we use a set of numerical experiments to derive insights into the performance and advantages of our RLR-CP solution framework in addressing such problems (Section 6.2).

6.1. Post-optimality analysis for modern prescriptive problems

The sensitivity and post-optimality analysis (SPA) of prescriptive models aims to understand how changes in model parameters affect the optimality of a recommended decision. Although this has been a longstanding topic of study in optimization and prescriptive analytics, the literature in the context of discrete optimization has been rather limited, with almost all work focusing on the classic sensitivity analysis of individual objective coefficients (see Section 2). By reframing SPA through the lens of inverse optimization, we broaden the set of questions that can be posed and enable richer insights into the solutions of modern optimization pipelines. To illustrate, we present examples in multi-objective optimization (Section 6.1.1) and contextual optimization (Section 6.1.2), both of which demand the analysis of parameters that underlie multiple objective coefficients.

6.1.1. Multi-objective optimization. Most real-world decision-making problems involve balancing multiple, potentially conflicting objectives. This is especially true in high-stakes domains and problems related to public policy. In practice, the predominant approach for solving multi-objective optimization problems is to aggregate objectives through a weighted-sum scalarization.

Specifically, consider a discrete optimization problem with d number of linear objectives, denoted as $\mathbf{f}_1^\top \mathbf{x}, \dots, \mathbf{f}_d^\top \mathbf{x}$. The weighted-sum scalarization of this problem can then be written as

$$\max_{\mathbf{x}} \left\{ \sum_{i=1}^d s_i \mathbf{f}_i^\top \mathbf{x} \mid \mathbf{x} \in \mathcal{X} \right\}, \quad (20)$$

where $s_i \geq 0$ denotes the weight assigned to that objective i . Without loss of generality, we can assume that $\sum_{i=1}^d s_i = 1$, where the value of s_i can be interpreted as the relative importance that is placed on the i -th objective.

Multi-objective discrete optimization problems of the form (20) are widely used to support decision-making in practice (e.g., Rönnqvist et al. 2017, Bertsimas et al. 2019, Chu et al. 2020, Peters et al. 2022, Christou et al. 2024, Ozel et al. 2025). Despite its widespread use, a key challenge that arises is the choice of the weights \mathbf{s} for (20). This task is complicated by the fact that different stakeholders often have different preferences over these weights, and reaching a consensus can become an extremely time consuming process (e.g., Kik et al. 2023, Umang et al. 2023).

In this context, it is valuable to understand and explore the range of weight vectors \mathbf{s} that are implicitly captured by a given solution, that is, the vectors \mathbf{s} that would all lead to the same solution

being generated by the prescriptive model. By definition, this requires extending the analysis of individual objective coefficients to the analysis of weight vectors that simultaneously influence multiple objective coefficients. We can formulate this generalized SPA problem as

$$\begin{aligned}
 & \min_{\boldsymbol{\theta}, \mathbf{w}} f(\mathbf{s}) \\
 & \text{s.t. } \boldsymbol{\theta} \in \mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X}) \\
 & \quad \boldsymbol{\theta} = \sum_i^d s_i \mathbf{f}_i \\
 & \quad \mathbf{s} \in \mathcal{S}
 \end{aligned} \tag{GSPA}$$

under various choices of $f(\mathbf{s})$ and \mathcal{S} . Our numerical experiments in Section 6.2 will be based on one such specification of (GSPA), which we introduce as an example below.

Example 4. Let $\hat{\mathbf{x}} \in \mathcal{X}$ be the optimal solution generated by (20) under a nominal set of weights $\hat{\mathbf{s}}$. Now, consider (GSPA) where $f(\mathbf{s}) = s_i$ and $\mathcal{S} = \{\mathbf{s} \mid \|\mathbf{s}\|_1 = 1, \mathbf{s} \geq \mathbf{0}\}$. This specification can be used to assess the flexibility in assigning a weight to a particular objective i . For instance, if s_i can be increased substantially compared to the nominal value \hat{s}_i , it implies that $\hat{\mathbf{x}}$ could still be generated under a much higher weighting of objective i , provided that the remaining weights are adjusted accordingly. Conversely, if s_i cannot vary by much, it suggest that a different prescriptive solution should be considered if one were to place more emphasis on objective i . More generally, we can use such specifications to explore the family of weights that ultimately yield the same solution.

6.1.2. Contextual optimization. Problems in the form of (GSPA) are also relevant in contextual and data-driven optimization, especially in predict-then-optimize pipelines. In this literature, we use a set of contextual attributes $\{\mathbf{f}_1, \dots, \mathbf{f}_m\} \in \mathbb{R}^n$ and a function $\hat{g}(\mathbf{f}_1, \dots, \mathbf{f}_m) \rightarrow \hat{\boldsymbol{\theta}}$ to generate a prediction for $\boldsymbol{\theta}$ in (FO). As an example, in standard routing problems aimed at minimizing total travel time, each feature $\mathbf{f}_i \in \mathbb{R}^d$ may encode information such as time-of-day, weather conditions, and traffic incidents. These contextual features are then used to predict the travel times along arcs of a road network, represented collectively by the parameter vector $\boldsymbol{\theta}$.

One of the most widely studied functions is the affine mapping, i.e.,

$$\hat{g}(\mathbf{f}_1, \dots, \mathbf{f}_m) := \beta_0 + \beta_1 \mathbf{f}_1 + \dots + \beta_d \mathbf{f}_m \rightarrow \hat{\boldsymbol{\theta}}, \tag{21}$$

where we denote $\boldsymbol{\beta} = [\beta_0, \dots, \beta_d]$ as the set of regression coefficients that are estimated using a data set (e.g., Elmachtoub and Grigas 2022).

Consider an optimal solution $\hat{\mathbf{x}}$ to $\text{FO}(\hat{\boldsymbol{\theta}}, \mathcal{X})$, where $\hat{\boldsymbol{\theta}}$ is generated via (21) for some set of feature vectors. Much like the multi-objective problem described in Section 6.1.1, the traditional single-coefficient sensitivity analysis of $\hat{\boldsymbol{\theta}}$ may be of limited value, since the objective coefficients are not

independent but are instead coupled through the features and regression parameters. Using variations of (GSPA), e.g., where \mathbf{s} is replaced with β , we can examine the sensitivity of a prescriptive decision to the values of β . Similarly, if we hold β fixed but define \mathbf{f} to be decision variables in the inverse problem, we could analyze the sensitivity to the contextual attributes \mathbf{f} .

6.2. Computational study

We present a computational study to explore the characteristics and performance of the RLR-CP Algorithm when solving SPA problems defined in the form of (GSPA). Our forward problems are defined to be multi-objective knapsack problems, which commonly arise in resource allocation and selection problems with competing objectives (e.g., Montoya-Weiss and Calantone 1999, Kulturel-Konak et al. 2023). The experimental setup is first outlined in Section 6.2.1, followed by key insights and discussion in Section 6.2.2.

6.2.1. Experimental setup. We consider multi-objective knapsack problems with 100, 200, 400, and 800 variables, each with 10 objective vectors. For each problem, we generate a solution $\hat{\mathbf{x}}$ using a nominal uniform weighting of $s_i = 0.1$ for each objective $i \in \{1, \dots, 10\}$. Then, we compute the maximum values that each value of s_i could take while $\hat{\mathbf{x}}$ remains optimal, by solving the problem described in Example 4. This leads to a total of 40 inverse discrete optimization problems spanning the four problem sizes.

For the RLR-CP Algorithm, we consider a static depth selection policy based on (19), where $\tilde{\eta} = \{0.0001, 0.001, 0.01, 0.1\}$. We adopt the trust-region based method from Bodur et al. (2022) as the CP-subroutine in our RLR-CP algorithm. Similarly, as benchmarks, we consider the standard CP algorithm (Wang 2009) and the trust-region based CP algorithm (Bodur et al. 2022). These simply correspond to versions of the RLR-CP Algorithm without the RLR component, i.e., where $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W})) = \mathbb{R}^n$ (see Section 5.2).

The full details of the instances, the specifications of the RLR-CP Algorithm, and the benchmarks are found in Section A4.2 of the Electronic Companion. All experiments are conducted on an Apple Mac M4 equipped with 10 cores and 10 threads and 16GBs of RAM using Gurobi 12.0.1 in Python.

6.2.2. Analysis and insights. We begin by summarizing the three key insights that are demonstrated through our computational experiments. The corresponding results and discussion that substantiate and elaborate on each insight is presented after the summary.

- *Insight I:* Under tight time constraints, the RLR-CP Algorithm can generate informative and high-quality solutions to (IO), whereas standalone CP algorithms are of little practical value.
- *Insight II:* As a general guideline, shorter time constraints should be paired with smaller depth parameters, as they tend to yield more frequent improvements in solution quality.

- *Insight III:* There exists a critical point in which further decreasing the depth parameter may degrade the performance of the RLR-CP Algorithm. This also suggests that the algorithm could further benefit from a more systematic tuning of the depth selection policy.

Tables 2 and 3 illustrates the aggregate results for solving the SPA problems over the instances with 200 and 400 variables, respectively; results for the other instances can be found in Section A4.3 in the Electronic Companion. Each table highlights the number of instances, out of 10, that are solved to a specified optimality gap within a given time limit.

Time Limit	25 seconds				50 seconds				100 seconds			
Optimality Gap	15%	5%	1%	0%	15%	5%	1%	0%	15%	5%	1%	0%
RLR-CP: $\tilde{\eta} = 0.0001$	3	0	0	0	9	5	1	0	10	10	9	4
RLR-CP: $\tilde{\eta} = 0.001$	9	7	1	1	9	9	9	9	10	10	10	10
RLR-CP: $\tilde{\eta} = 0.01$	0	0	0	0	3	3	3	3	9	9	9	9
RLR-CP: $\tilde{\eta} = 0.1$	0	0	0	0	0	0	0	0	5	5	5	5
CP: Standard	0	0	0	0	0	0	0	0	6	6	6	6
CP: Trust Region	0	0	0	0	3	3	3	3	9	9	9	9

Table 2 Number of instances reaching specified optimality gaps within time limits, for 200 variable instances.

Time Limit	200 seconds				500 seconds				1000 seconds			
Optimality Gap	15%	5%	1%	0%	15%	5%	1%	0%	15%	5%	1%	0%
RLR-CP: $\tilde{\eta} = 0.0001$	10	9	3	0	10	10	10	3	10	10	10	9
RLR-CP: $\tilde{\eta} = 0.001$	3	2	1	0	9	8	7	7	10	10	10	10
RLR-CP: $\tilde{\eta} = 0.01$	0	0	0	0	1	1	1	1	8	8	8	8
RLR-CP: $\tilde{\eta} = 0.1$	0	0	0	0	0	0	0	0	5	5	5	5
CP: Standard	0	0	0	0	0	0	0	0	6	6	6	6
CP: Trust Region	0	0	0	0	0	0	0	0	7	7	7	7

Table 3 Number of instances reaching specified optimality gaps within time limits, for 400 variable instances.

We first use Tables 2 and 3 to examine the performance the RLR-CP Algorithm relative to standalone CP algorithms. Recall that CP algorithms do not generate (or improve) inverse-feasible solutions until the final iteration, at which point it generates the optimal one. Because of this property, CP algorithms can be entirely uninformative under time constraints. For example, when we consider a time limit of 25 seconds in Table 2 or 500 seconds in Table 3, CP algorithms do not generate any inverse-feasible solution within any level of optimality. In contrast, the RLR-CP Algorithm with smaller values of $\tilde{\eta}$ can generate near-optimal solutions within the same allotted time. For example, in 500 seconds, all instances in Table 3 could be solved to 1% optimality for the

smallest value of $\tilde{\eta}$. Furthermore, the RLR-CP Algorithm can progressively refine these inverse-feasible solutions over time; by 1000 seconds, nearly all instances across all values of $\tilde{\eta}$ are solved to optimality. These results demonstrate that the RLR-CP Algorithm is effective both at efficiently generating high-quality solutions and at solving inverse optimization problems to optimality.

The second insight that can be drawn from these tables is that increasing the depth parameter can potentially lead to a significant degradation in performance. This pattern is consistent with both the theoretical results and our expectations: as the value of $\tilde{\eta}$ increases, the RLR-CP Algorithm behaves more like the standalone CP algorithm, which suffers from the limitations discussed in the previous paragraph (and also in Section 5.2). This is reflected in the results of $\tilde{\eta} = 0.1$. Conversely, when we decrease the depth parameter too much beyond a certain point, we also diminish the ability to make significant improvements. This is reflected in the results of $\tilde{\eta} = 0.0001$ in Table 2.

Figure 6 provides a more detailed illustration on the performance of each algorithmic specification for a particular instance of the 200-variable model. In this instance, we seek to maximize the value of the weight of the 7-th objective vector, s_7 , in (GSPA). In the optimal solution \mathbf{s}^* , $s_7^* = 0.0764$. The figure shows a plot of the values of \tilde{s}_7^t in all candidate vectors $\tilde{\mathbf{s}}^t$ that are generated in each iteration of the cutting plane subroutine, as well as those of the standalone CP algorithm.

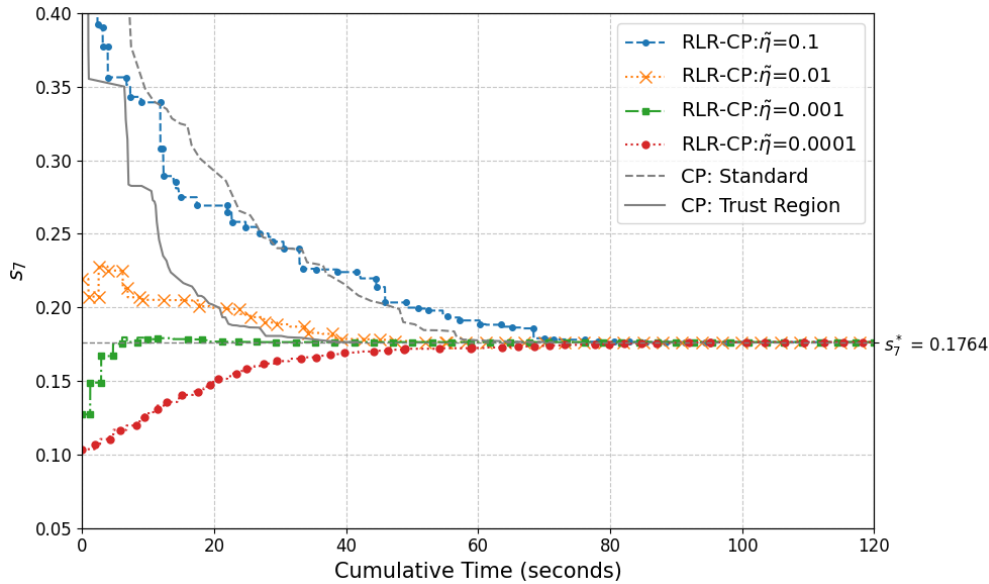


Figure 6 The performance of different algorithms for a single 200 variable instance.

To best interpret this figure, we first note that all values above s_7^* , by definition, correspond to candidate vectors \mathbf{s}^t that are not inverse-feasible (otherwise, s_7^* cannot be optimal). On the other hand, for the line segments that lie below s_7^* , any point on the line that precedes a shift towards s_7^* must correspond to an inverse-feasible vector (as the increase can only result from adding a new

restricting inequality using a new inverse-feasible vector). With this in mind, we first remark that both CP algorithms, as expected, converge towards the value of s_7^* from above. This is also true for the RLR-CP Algorithm for larger values of $\tilde{\eta}$, i.e., when $\tilde{\eta} = 0.1$ and $\tilde{\eta} = 0.01$. However, note that when we further decrease $\tilde{\eta}$ to 0.001, the algorithm initially generates a few inverse-feasible vectors, followed by a few cuts, and then quickly converges to s_7^* . Finally, when we further decrease $\tilde{\eta}$ to 0.0001, the algorithm continues to generate and improve inverse-feasible vectors, albeit at a slower rate. These observations also demonstrate why it can be useful to view the RLR-CP Algorithm as a more complete primal-dual solution framework for inverse optimization, where the depth parameter implicitly controls the time spent improving dual bounds versus primal (i.e., inverse-feasible) solutions.

7. Conclusion

Inverse optimization has emerged as a powerful framework for analyzing model parameters that rationalize observed or prescribed decisions. In this paper, we take a holistic approach to solving and applying IO models for prescriptive problems that have discrete decision variables, and make three main contributions. We first introduce a versatile and novel technique for generating rich approximations of the inverse-feasible region using the concept of restricting linear relaxations. We then leverage RLRs to develop a generalized solution framework for inverse optimization, and show that this framework directly addresses the main limitation of standalone cutting plane approaches. Finally, we use inverse optimization to develop a modern perspective on sensitivity and post-optimality analysis, through which we derive insights into the models and algorithms proposed in the paper. We conclude by providing two general directions for future research.

Direct extensions. The primary goal our paper is to propose a family of new ideas and methods for inverse optimization that are broadly applicable across discrete optimization problems. Nevertheless, it is natural to expect that these ideas can potentially be further refined by focusing on specific classes of forward problems, for example those with common underlying structures (e.g., knapsack, matching, packing, or routing problems). In such settings, it may be possible to derive stronger theoretical results, or integrate existing techniques – such as problem-specific valid inequalities – into the methodological framework that we have developed.

Broader applications. New applications of inverse optimization and the inverse-feasible region continue to emerge. Exploring problems and domains beyond those studied in this paper may, in turn, motivate new methodological questions that require extending the approaches proposed here. For example, Section 2.3 outlines several applications that use the inverse-feasible region but that do not precisely fit within the form of the IO problems studied in this paper. These examples were included to highlight the broader potential of the techniques developed here, and we believe there are many promising avenues for extending or developing new ideas in this space.

References

- Agarwal R, Ergun Ö (2010) Network design and allocation mechanisms for carrier alliances in liner shipping. *Operations Research* 58(6):1726–1742.
- Ahuja RK, Orlin JB (2001) Inverse optimization. *Operations Research* 49(5):771–783.
- Andersen KA, Boomsma TK, Nielsen LR (2023) MILP sensitivity analysis for the objective function coefficients. *INFORMS Journal on Optimization* 5(1):92–109.
- Aswani A, Shen ZJ, Siddiq A (2018) Inverse optimization with noisy data. *Operations Research* 66(3):870–892.
- Berden S, Mahmutogullari AI, Tsouros D, Guns T (2025) Solver-free decision-focused learning for linear optimization problems. *arXiv preprint arXiv:2505.22224* .
- Bertsimas D, Delarue A, Martin S (2019) Optimizing schools’ start time and bus routes. *Proceedings of the National Academy of Sciences* 116(13):5943–5948.
- Bertsimas D, Tsitsiklis JN (1997) *Introduction to linear optimization*, volume 6 (Athena scientific Belmont, MA).
- Besbes O, Fonseca Y, Lobel I (2025) Contextual inverse optimization: Offline and online learning. *Operations Research* 73(1):424–443.
- Bodur M, Chan TC, Zhu IY (2022) Inverse mixed integer optimization: Polyhedral insights and trust region methods. *INFORMS Journal on Computing* 34(3):1471–1488.
- Bulut A, Ralphs TK (2021) On the complexity of inverse mixed integer linear optimization. *SIAM Journal on Optimization* 31(4):3014–3043.
- Chan T, Delage E, Lin B (2024) Conformal inverse optimization for adherence-aware prescriptive analytics. *Available at SSRN* .
- Chan TC, Mahmood R, Zhu IY (2025) Inverse optimization: Theory and applications. *Operations Research* 73(2):1046–1074.
- Christou IT, Vagianou E, Vardoulas G (2024) Planning courses for student success at the american college of greece. *INFORMS Journal on Applied Analytics* 54(4):365–379.
- Chu A, Keskinocak P, Villarreal MC (2020) Empowering denver public schools to optimize school bus operations. *INFORMS Journal on Applied Analytics* 50(5):298–312.
- Chvátal V, Cook W, Espinoza D (2013) Local cuts for mixed-integer programming. *Mathematical Programming Computation* 5(2):171–200.
- Cornuéjols G (2008) Valid inequalities for mixed integer linear programs. *Mathematical Programming* 112(1):3–44.
- Duan Z, Wang L (2011) Heuristic algorithms for the inverse mixed integer linear programming problem. *Journal of Global Optimization* 51:463–471.

- D'Ambrosio C, Frangioni A, Liberti L, Lodi A (2010) On interval-subgradient and no-good cuts. *Operations Research Letters* 38(5):341–345.
- Elmachtoub AN, Grigas P (2022) Smart “predict, then optimize”. *Management Science* 68(1):9–26.
- Geoffrion AM, Nauss R (1977) Parametric and postoptimality analysis in integer linear programming. *Management Science* 23(5):453–466.
- Heuberger C (2004) Inverse combinatorial optimization: a survey on problems, methods, and results. *Journal of Combinatorial Optimization* 8:329–361.
- Jenkins L (1982) Parametric mixed integer programming: an application to solid waste management. *Management Science* 28(11):1270–1284.
- Kik D, Wichmann MG, Spengler TS (2023) Small-or medium-sized enterprise uses operations research to select and develop its headquarters location. *INFORMS Journal on Applied Analytics* 53(4):312–331.
- Kılınç-Karzan F, Toriello A, Ahmed S, Nemhauser G, Savelsbergh M (2009) Approximating the stability region for binary mixed-integer programs. *Operations Research Letters* 37(4):250–254.
- Klein D, Holm S (1979) Integer programming post-optimal analysis with cutting planes. *Management Science* 25(1):64–72.
- Kulturel-Konak S, Konak A, Jakielaszek L, Gavirneni N (2023) Menu engineering for continuing care senior living facilities with captive dining patrons. *INFORMS Journal on Applied Analytics* 53(3):218–239.
- Lamperski JB, Schaefer AJ (2015) A polyhedral characterization of the inverse-feasible region of a mixed-integer program. *Operations Research Letters* 43(6):575–578.
- Lin B, Delage E, Chan T (2024) Conformal inverse optimization. *Advances in Neural Information Processing Systems* 37:63534–63564.
- Liu L, Qi X, Xu Z (2024) Stabilizing grand cooperation via cost adjustment: An inverse optimization approach. *INFORMS Journal on Computing* 36(2):635–656.
- Marcotte P, Mercier A, Savard G, Verter V (2009) Toll policies for mitigating hazardous materials transport risk. *Transportation Science* 43(2):228–243.
- Mohajerin Esfahani P, Shafieezadeh-Abadeh S, Hanasusanto GA, Kuhn D (2018) Data-driven inverse optimization with imperfect information. *Mathematical Programming* 167(1):191–234.
- Montoya-Weiss M, Calantone RJ (1999) Development and implementation of a segment selection procedure for industrial product markets. *Marketing Science* 18(3):373–395.
- Niendorf M, Girard AR (2017) Exact and approximate stability of solutions to traveling salesman problems. *IEEE Transactions on Cybernetics* 48(2):583–595.
- Niendorf M, Kabamba PT, Girard AR (2015) Stability of solutions to classes of traveling salesman problems. *IEEE Transactions on Cybernetics* 46(4):973–985.

- Ozel A, Smilowitz K, Goldstein LK (2025) Community-engaged school district design: A stream-based approach. *Operations Research* (Articles in Advance).
- Peters K, Silva S, Wolter TS, Anjos L, van Ettekovén N, Combette É, Melchiori A, Fleuren H, den Hertog D, Ergun Ö (2022) UN world food programme: toward zero hunger with analytics. *INFORMS journal on applied analytics* 52(1):8–26.
- Petersen CC (1967) Computational experience with variants of the Balas algorithm applied to the selection of R&D projects. *Management Science* 13(9):736–750.
- Rönnqvist M, Svenson G, Flisberg P, Jönsson LE (2017) Calibrated route finder: Improving the safety, environmental consciousness, and cost effectiveness of truck routing in sweden. *Interfaces* 47(5):372–395.
- Roodman GM (1974) Postoptimality analysis in integer programming by implicit enumeration: The mixed integer case. *Naval Research Logistics Quarterly* 21(4):595–607.
- Schaefer AJ (2009) Inverse integer programming. *Optimization Letters* 3:483–489.
- Schrage L, Wolsey L (1985) Sensitivity analysis for branch and bound integer programming. *Operations Research* 33(5):1008–1023.
- Shahmoradi Z, Lee T (2022) Quantile inverse optimization: Improving stability in inverse linear programming. *Operations research* 70(4):2538–2562.
- Tang B, Khalil EB (2024) Cave: A cone-aligned approach for fast predict-then-optimize with binary linear programs. *International Conference on the Integration of Constraint Programming, Artificial Intelligence, and Operations Research*, 193–210 (Springer).
- Tavashoğlu O, Lee T, Valeva S, Schaefer AJ (2018) On the structure of the inverse-feasible region of a linear program. *Operations Research Letters* 46(1):147–152.
- Umang N, Kiefer D, Gonzalez P, Thevenot H, Johnson CD, Martin J, Stephenson E, Cline J, Gemici-Ozkan B, Beniaminy I (2023) Practice summary: GE optimizes for aircraft engine overhaul scheduling and shop assignment. *INFORMS Journal on Applied Analytics* 53(2):128–132.
- Wang L (2009) Cutting plane algorithms for the inverse mixed integer linear programming problem. *Operations Research Letters* 37(2):114–116.
- Wei N, Zhang P (2024) Adjustability in robust linear optimization. *Mathematical Programming* 208(1):581–628.
- Zattoni Scroccaro P, Atasoy B, Mohajerin Esfahani P (2025) Learning in inverse optimization: Incenter cost, augmented suboptimality loss, and algorithms. *Operations Research* 73(5):2661–2679.
- Zhou W, Orfanoudaki A, Zhu S (2025) Conformalized decision risk assessment. *arXiv preprint arXiv:2505.13243* .

Electronic Companion to “Inverse Optimization with Discrete Decisions”

Section A1 contains the proofs of all results presented in the main body of the paper. Sections A2, A3 and A4 contain additional materials for Sections 4, 5 and 6, respectively.

A1. Proofs

Proofs are presented in the same order as the results in the paper.

A1.1. Proof of Proposition 1

We can reformulate the cone $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}))$ as follows, starting with its original definition:

$$\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W})) = \{\boldsymbol{\theta} \mid \boldsymbol{\theta}^\top \hat{\mathbf{x}} \geq \boldsymbol{\theta}^\top \mathbf{x}, \quad \forall \mathbf{x} \in \mathcal{X}^{\text{LR}}\} \quad (\text{A1a})$$

$$= \{\boldsymbol{\theta} \mid \boldsymbol{\theta}^\top \hat{\mathbf{x}} \geq \max\{\boldsymbol{\theta}^\top \mathbf{x} \mid \forall \mathbf{x} \in \mathcal{X}^{\text{LR}}\}\} \quad (\text{A1b})$$

$$= \{\boldsymbol{\theta} \mid \boldsymbol{\theta}^\top \hat{\mathbf{x}} \geq \max\{\boldsymbol{\theta}^\top \mathbf{x} \mid \mathbf{A}\mathbf{x} \leq \mathbf{b}, \mathbf{W}\mathbf{x} \leq \boldsymbol{\tau}\}\} \quad (\text{A1c})$$

$$= \{\boldsymbol{\theta} \mid \boldsymbol{\theta}^\top \hat{\mathbf{x}} \geq \min\{\mathbf{p}^\top \mathbf{b} + \mathbf{z}^\top \boldsymbol{\tau} \mid \mathbf{A}^\top \mathbf{p} + \mathbf{W}^\top \mathbf{z} = \boldsymbol{\theta}, \mathbf{p} \geq \mathbf{0}, \mathbf{z} \geq \mathbf{0}\}\} \quad (\text{A1d})$$

$$= \{\boldsymbol{\theta} \mid \boldsymbol{\theta}^\top \hat{\mathbf{x}} \geq \mathbf{p}^\top \mathbf{b} + \mathbf{z}^\top \boldsymbol{\tau}, \mathbf{A}^\top \mathbf{p} + \mathbf{W}^\top \mathbf{z} = \boldsymbol{\theta}, \mathbf{p} \geq \mathbf{0}, \mathbf{z} \geq \mathbf{0}\} \quad (\text{A1e})$$

The reformulation from (A1c) to (A1d) uses strong duality, where \mathbf{p} and \mathbf{z} are the dual variables associated with constraints $\mathbf{A}\mathbf{x} \leq \mathbf{b}$ and $\mathbf{W}\mathbf{x} \leq \boldsymbol{\tau}$, respectively.

A1.2. Proof of Theorem 1

1. Proof that $\mathbf{w} \in \text{int}(\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W})))$ for all $\mathbf{w} \in \mathbf{W}$: Note that since $\mathcal{X}^{\text{LR}}(\mathcal{W})$ is defined by restricting inequalities $\mathbf{w}^\top \mathbf{x} \leq \tau$ where $\mathbf{w}^\top \hat{\mathbf{x}} > \tau$, it must be true that $\mathbf{w}^\top \hat{\mathbf{x}} > \mathbf{w}^\top \mathbf{x}$ must hold for all $\mathbf{x} \in \mathcal{X}^{\text{LR}}(\mathcal{W})$. Since there are no binding constraints at \mathbf{w} , this implies that $\mathbf{w} \in \text{int}(\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W})))$.
2. Proof that $\dim(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W})) = n$: Since there exists a $\mathbf{w} \in \text{int}(\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W})))$, this also implies that $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}})$ must be full-dimensional, i.e., $\dim(\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}})) = n$.
3. Proof that $\dim(\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W})) \cap \mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})) = \dim(\mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X}))$: Suppose $\dim(\mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})) = m$ where $m \leq n$. This implies that we can generate m number of linearly independent vectors $\boldsymbol{\theta}^1, \dots, \boldsymbol{\theta}^m \in \mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$. Consider any $\mathbf{w} \in \mathbf{W}$. Since $\mathbf{w} \in \mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$ and $\mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$ is convex, any convex combination of \mathbf{w} and $\boldsymbol{\theta}^i$ where $i = \{1, \dots, m\}$ must also be inverse-feasible. Furthermore, note that since $\mathbf{w} \in \text{int}(\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W})))$, there must exist a $\mathcal{B}_\epsilon(\mathbf{w}) \subseteq \mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}))$ where $\mathcal{B}_\epsilon(\mathbf{w}) := \{\mathbf{w} + \boldsymbol{\delta} \mid \|\boldsymbol{\delta}\|_2 \leq \epsilon\}$ for some $\epsilon > 0$. Thus, for any vector $(\boldsymbol{\theta}^i - \mathbf{w})$, there must exist a sufficiently small $0 < \lambda^i \leq 1$ for which $\mathbf{w} + \lambda^i(\boldsymbol{\theta}^i - \mathbf{w})$ lies in $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}))$; note that this $\mathbf{w} + \lambda^i(\boldsymbol{\theta}^i - \mathbf{w})$ also belongs to $\mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$, as it is a convex combination of \mathbf{w} and $\boldsymbol{\theta}^i$. Finally, since

$$\text{span}(\mathbf{w}, \mathbf{w} + \lambda^1(\boldsymbol{\theta}^1 - \mathbf{w}), \dots, \mathbf{w} + \lambda^m(\boldsymbol{\theta}^m - \mathbf{w})) = m, \quad (\text{A2})$$

it must be true that $\dim(\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W})) \cap \mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})) = m$. □

A1.3. Proof of Theorem 2

(\Leftarrow) We first prove that if $\mathcal{X}^{\text{LR}}(\mathcal{W})$ is a pure-RLR, then $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W})) \subseteq \mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$. The main idea of this proof is to leverage the fact that $\mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$ is equivalent to $\mathcal{K}(\hat{\mathbf{x}}, \text{ext}(\mathcal{X}) \setminus \{\hat{\mathbf{x}}\})$, i.e.,

$$\begin{aligned} \mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X}) &= \{\boldsymbol{\theta} \mid \boldsymbol{\theta}^\top \hat{\mathbf{x}} \geq \boldsymbol{\theta}^\top \mathbf{x} \quad \forall \mathbf{x} \in \mathcal{X}\} \\ &= \{\boldsymbol{\theta} \mid \boldsymbol{\theta}^\top \hat{\mathbf{x}} \geq \boldsymbol{\theta}^\top \mathbf{x} \quad \forall \mathbf{x} \in \text{conv}(\mathcal{X})\} \\ &= \{\boldsymbol{\theta} \mid \boldsymbol{\theta}^\top \hat{\mathbf{x}} \geq \boldsymbol{\theta}^\top \mathbf{x} \quad \forall \mathbf{x} \in \text{ext}(\mathcal{X})\} \\ &= \{\boldsymbol{\theta} \mid \boldsymbol{\theta}^\top \hat{\mathbf{x}} \geq \boldsymbol{\theta}^\top \mathbf{x} \quad \forall \mathbf{x} \in \text{ext}(\mathcal{X}) \setminus \{\hat{\mathbf{x}}\}\} \\ &= \mathcal{K}(\hat{\mathbf{x}}, \text{ext}(\mathcal{X}) \setminus \{\hat{\mathbf{x}}\}). \end{aligned}$$

Let $\tilde{\mathcal{X}} := \text{ext}(\mathcal{X}) \setminus \{\hat{\mathbf{x}}\}$. Since $\mathcal{X}^{\text{LR}}(\mathcal{W})$ is a pure-RLR, $\tilde{\mathcal{X}} \subset \mathcal{X}^{\text{LR}}(\mathcal{W})$ by definition. Thus, it must be true that $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W})) \subseteq \mathcal{K}(\hat{\mathbf{x}}, \tilde{\mathcal{X}}) = \mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$, as $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}') \subseteq \mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}'')$ for any $\mathcal{X}'' \subseteq \mathcal{X}'$.

(\Rightarrow) We now show that if $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W})) \subseteq \mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$ then $\mathcal{X}^{\text{LR}}(\mathcal{W})$ must be a pure-RLR. We consider a proof by contradiction, namely, that there exists a mixed-RLR $\mathcal{X}^{\text{LR}}(\mathcal{W})$ for which $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W})) \subseteq \mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$.

In order to prove this statement, we first note that, by Theorem 1, since $\dim(\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}))) = n$, it must be true that $\mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$ and that $\hat{\mathbf{x}} \in \text{ext}(\mathcal{X})$. Furthermore, since $\mathcal{X}^{\text{LR}}(\mathcal{W})$ is a mixed-RLR, it must be true that there exists a $\tilde{\mathbf{x}} \in \text{ext}(\mathcal{X}) \setminus \{\hat{\mathbf{x}}\}$ such that $\tilde{\mathbf{x}} \notin \mathcal{X}^{\text{LR}}(\mathcal{W})$.

The proof follows in three parts:

Part 1 We first show that if there exists a $\tilde{\mathbf{x}} \in \text{ext}(\mathcal{X}) \setminus \{\hat{\mathbf{x}}\}$ where $\mathbf{w}^\top \tilde{\mathbf{x}} > \tau$, then there must also exist a $\bar{\mathbf{x}} \in \text{ext}(\mathcal{X}) \setminus \{\hat{\mathbf{x}}\}$ that is adjacent to $\hat{\mathbf{x}}$ that also satisfies $\mathbf{w}^\top \bar{\mathbf{x}} > \tau$.

Part 2 We then show that, for this $\bar{\mathbf{x}} \in \text{ext}(\mathcal{X}) \setminus \{\hat{\mathbf{x}}\}$, there exists a $\boldsymbol{\theta}$ such that $\bar{\mathbf{x}}$ is the only solution to $\max\{\boldsymbol{\theta}^\top \mathbf{x} \mid \mathbf{x} \in \text{ext}(\mathcal{X}) \setminus \{\hat{\mathbf{x}}\}\}$ where $\boldsymbol{\theta}^\top \hat{\mathbf{x}} = \boldsymbol{\theta}^\top \bar{\mathbf{x}}$. We show that this $\boldsymbol{\theta}$, by definition, must be on the boundary of $\mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$.

Part 3 Finally, we show that this particular $\boldsymbol{\theta}$ is in the interior of $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}))$. This implies that $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W})) \not\subseteq \mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$.

Proofs of Parts 1 and 2: To avoid repetition, the proofs of Parts 1 and 2 follow directly from the analogous arguments in Parts 1 and 2 of Section A1.5, by replacing \mathcal{X}^{LR} in A1.5 with $\text{conv}(\mathcal{X})$, since $\text{ext}(\mathcal{X})$ is equivalent to $\text{ext}(\text{conv}(\mathcal{X}))$ (Section 1.2).

Proof of Part 3: We first show that there cannot exist any $\mathbf{x} \in \mathcal{X}^{\text{LR}}(\mathcal{W})$ for which $\hat{\boldsymbol{\theta}}^\top \hat{\mathbf{x}} = \hat{\boldsymbol{\theta}}^\top \mathbf{x}$. To do so, we define the line segment between $\hat{\mathbf{x}}$ and $\bar{\mathbf{x}}$ as $\mathcal{L} := \{\mathbf{y} \mid \mathbf{y} = \rho \hat{\mathbf{x}} + (1 - \rho)(\bar{\mathbf{x}}), \rho \in [0, 1]\}$. Note that by the definition of $\hat{\boldsymbol{\theta}}$, $\mathcal{L} \subseteq \text{conv}(\mathcal{X})$ describes the complete set of solutions $\mathbf{x} \in \text{conv}(\mathcal{X})$ for which $\hat{\boldsymbol{\theta}}^\top \hat{\mathbf{x}} = \hat{\boldsymbol{\theta}}^\top \mathbf{x}$. In other words, $\hat{\boldsymbol{\theta}}^\top \mathbf{y} < \hat{\boldsymbol{\theta}}^\top \mathbf{x} \quad \forall \mathbf{y} \in \text{conv}(\mathcal{X}) \setminus \mathcal{L}$.

We now show that $\mathcal{L} \not\subseteq \mathcal{X}^{\text{LR}}(\mathcal{W})$. This is straightforward, as $\mathbf{w}^\top \bar{\mathbf{x}} > \tau$ and $\mathbf{w}^\top \hat{\mathbf{x}} > \tau$, which implies that $\mathbf{w}^\top \mathbf{y} > \tau$ for all $\mathbf{y} \in \mathcal{L}$. This implies that $\hat{\boldsymbol{\theta}}^\top (\hat{\mathbf{x}} - \mathbf{x}) > 0 \quad \forall \mathbf{x} \in \mathcal{X}^{\text{LR}}(\mathcal{W})$, which by definition implies that $\hat{\boldsymbol{\theta}}$ is in the interior of $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}))$. Since $\boldsymbol{\theta}$ is on the boundary of $\mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$ but is in the interior of $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}))$, it must be true that $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W})) \not\subseteq \mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$.

This contradicts the statement that $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W})) \subseteq \mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$, which implies that $\mathcal{X}^{\text{LR}}(\mathcal{W})$ cannot be a mixed-RLR. \square

A1.4. Proof of Proposition 2

First, note that constraint (14c) excludes only the solution $\hat{\mathbf{x}} \in \mathcal{X}$ from \mathcal{X} , as it is a reformulation of the no-good cut $|\mathbf{x} - \hat{\mathbf{x}}|_1 \geq 1$ for binary programs (D'Ambrosio et al. 2010). Second, since $\mathcal{X} = \text{ext}(\mathcal{X})$ when $\mathcal{X} \subseteq \{0, 1\}^n$, solving Problem (14) is equivalent to finding the maximizer of $\mathbf{w}^\top \mathbf{x}$ where $\mathbf{x} \in \text{ext}(\mathcal{X}) \setminus \hat{\mathbf{x}}$. Let $\tilde{\mathbf{x}}$ denote this maximizer, i.e., $\tau^*(\mathbf{w}) := \mathbf{w}^\top \tilde{\mathbf{x}}$. If $\tau^*(\mathbf{w}) < \mathbf{w}^\top \hat{\mathbf{x}}$, then $\hat{\mathbf{x}} \in \mathcal{X}$ must be the unique optimal solution under \mathbf{w} , and $\mathbf{w}^\top \mathbf{x} \leq \tau^*(\mathbf{w})$ would give a tight separating inequality. Otherwise, $\mathbf{w}^\top \mathbf{x} \leq \tau^*(\mathbf{w})$ cannot form a separating hyperplane as $\tau^*(\mathbf{w}) \geq \mathbf{w}^\top \hat{\mathbf{x}}$; instead, it generates a valid inequality of the linear relaxation. This concludes the proof. \square

A1.5. Proof of Theorem 3

(\Rightarrow) First, we prove that if $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}) \subset \mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}))$, then there exists a $\mathbf{x} \in \text{ext}(\mathcal{X}^{\text{LR}}) \setminus \{\hat{\mathbf{x}}\}$ where $\mathbf{w}^\top \mathbf{x} > \tau$. We prove this by contradiction. Specifically, suppose that $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}) \subset \mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}))$, but that $\text{ext}(\mathcal{X}^{\text{LR}}) \setminus \{\hat{\mathbf{x}}\} = \text{ext}(\mathcal{X}^{\text{LR}}(\mathcal{W})) \setminus \{\hat{\mathbf{x}}\}$. Through the following set of reformulations, we prove that $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}) = \mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}))$, which is a contradiction:

$$\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}) = \{\boldsymbol{\theta} \mid \boldsymbol{\theta}^\top \hat{\mathbf{x}} \geq \boldsymbol{\theta}^\top \mathbf{x}, \forall \mathbf{x} \in \mathcal{X}^{\text{LR}}\} \quad (\text{A3a})$$

$$= \{\boldsymbol{\theta} \mid \boldsymbol{\theta}^\top \hat{\mathbf{x}} \geq \boldsymbol{\theta}^\top \mathbf{x}, \forall \mathbf{x} \in \text{ext}(\mathcal{X}^{\text{LR}})\} \quad (\text{A3b})$$

$$= \{\boldsymbol{\theta} \mid \boldsymbol{\theta}^\top \hat{\mathbf{x}} \geq \boldsymbol{\theta}^\top \mathbf{x}, \forall \mathbf{x} \in \text{ext}(\mathcal{X}^{\text{LR}}) \setminus \{\hat{\mathbf{x}}\}\} \quad (\text{A3c})$$

$$= \{\boldsymbol{\theta} \mid \boldsymbol{\theta}^\top \hat{\mathbf{x}} \geq \boldsymbol{\theta}^\top \mathbf{x}, \forall \mathbf{x} \in \text{ext}(\mathcal{X}^{\text{LR}}(\mathcal{W})) \setminus \{\hat{\mathbf{x}}\}\} \quad (\text{A3d})$$

$$= \{\boldsymbol{\theta} \mid \boldsymbol{\theta}^\top \hat{\mathbf{x}} \geq \boldsymbol{\theta}^\top \mathbf{x}, \forall \mathbf{x} \in \text{ext}(\mathcal{X}^{\text{LR}}(\mathcal{W}))\} \quad (\text{A3e})$$

$$= \{\boldsymbol{\theta} \mid \boldsymbol{\theta}^\top \hat{\mathbf{x}} \geq \boldsymbol{\theta}^\top \mathbf{x}, \forall \mathbf{x} \in \mathcal{X}^{\text{LR}}(\mathcal{W})\} \quad (\text{A3f})$$

$$= \mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W})). \quad (\text{A3g})$$

Equation (A3b) follows from (A3a) by noting that (A3a) can be equivalently rewritten as

$$\{\boldsymbol{\theta} \mid \boldsymbol{\theta}^\top \hat{\mathbf{x}} \geq \max\{\boldsymbol{\theta}^\top \mathbf{x} \mid \mathbf{x} \in \mathcal{X}^{\text{LR}}\}\}.$$

Since $\max\{\boldsymbol{\theta}^\top \mathbf{x} \mid \mathbf{x} \in \mathcal{X}^{\text{LR}}\}$ is a linear program, this is also equivalent to

$$\{\boldsymbol{\theta} \mid \boldsymbol{\theta}^\top \hat{\mathbf{x}} \geq \max\{\boldsymbol{\theta}^\top \mathbf{x} \mid \mathbf{x} \in \text{ext}(\mathcal{X}^{\text{LR}})\}\},$$

from which equation (A3b) is derived.

Equation (A3c) holds since $\boldsymbol{\theta}^\top \hat{\mathbf{x}} \geq \boldsymbol{\theta}^\top \hat{\mathbf{x}}$ is always satisfied. Equation (A3d) holds from the assumption that $\text{ext}(\mathcal{X}^{\text{LR}}) \setminus \{\hat{\mathbf{x}}\} = \text{ext}(\mathcal{X}^{\text{LR}}(\mathcal{W})) \setminus \{\hat{\mathbf{x}}\}$, i.e., the assumption we made to set up the proof by contradiction. Finally, equation (A3e) and (A3f) follow from the same line of reasoning used for (A3c) and (A3b), respectively. Since we have proved that $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}) = \mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}))$, this contradicts our assumption, and the original statement must be true.

(\Leftarrow) We now show that if there exists a $\tilde{\mathbf{x}} \in \text{ext}(\mathcal{X}^{\text{LR}}) \setminus \{\hat{\mathbf{x}}\}$ where $\mathbf{w}^\top \tilde{\mathbf{x}} > \tau$, then it must be true that $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}) \subset \mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}))$. Our proof focuses on the case where $\hat{\mathbf{x}} \in \text{ext}(\mathcal{X}^{\text{LR}})$. The case where $\hat{\mathbf{x}} \notin \text{ext}(\mathcal{X}^{\text{LR}})$ is easier to prove: if $\hat{\mathbf{x}} \notin \text{ext}(\mathcal{X}^{\text{LR}})$, then for any inverse-feasible vector \mathbf{w} , there must exist $\tilde{\mathbf{x}} \in \text{ext}(\mathcal{X}^{\text{LR}})$ such that $\mathbf{w}^\top \tilde{\mathbf{x}} \geq \mathbf{w}^\top \hat{\mathbf{x}}$. Thus, this would be an example of a $\tilde{\mathbf{x}} \in \text{ext}(\mathcal{X}^{\text{LR}}) \setminus \{\hat{\mathbf{x}}\}$ where $\mathbf{w}^\top \tilde{\mathbf{x}} > \tau$. Furthermore, since $\dim(\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}})) < n$ when $\hat{\mathbf{x}} \notin \text{ext}(\mathcal{X}^{\text{LR}})$, and since $\dim(\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}))) = n$, it must be true that $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}) \subset \mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}))$.

Thus, we focus on the case where $\hat{\mathbf{x}} \in \text{ext}(\mathcal{X}^{\text{LR}})$, and our proof unfolds in three parts:

Part 1 We first show that if there exists a $\tilde{\mathbf{x}} \in \text{ext}(\mathcal{X}^{\text{LR}}) \setminus \{\hat{\mathbf{x}}\}$ where $\mathbf{w}^\top \tilde{\mathbf{x}} > \tau$, then there must also exist a $\bar{\mathbf{x}} \in \text{ext}(\mathcal{X}^{\text{LR}}) \setminus \{\hat{\mathbf{x}}\}$ that is adjacent to $\hat{\mathbf{x}}$ that also satisfies $\mathbf{w}^\top \bar{\mathbf{x}} > \tau$.

Part 2 We then show that, for this $\bar{\mathbf{x}} \in \text{ext}(\mathcal{X}^{\text{LR}}) \setminus \{\hat{\mathbf{x}}\}$, there exists a $\boldsymbol{\theta}$ such that $\bar{\mathbf{x}}$ is the only solution to $\max\{\boldsymbol{\theta}^\top \mathbf{x} \mid \mathbf{x} \in \text{ext}(\mathcal{X}^{\text{LR}}) \setminus \{\hat{\mathbf{x}}\}\}$ where $\boldsymbol{\theta}^\top \hat{\mathbf{x}} = \boldsymbol{\theta}^\top \bar{\mathbf{x}}$. We show that this $\boldsymbol{\theta}$, by definition, must be on the boundary of $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}})$.

Part 3 Finally, we show that this particular $\boldsymbol{\theta}$ is in the interior of $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}))$. This implies that $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}) \subset \mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}))$.

We now provide the proof for each of these parts.

Proof of Part 1: Let $\widehat{\text{ext}}(\mathcal{X}^{\text{LR}})$ denote the extreme points of \mathcal{X}^{LR} that are adjacent to $\hat{\mathbf{x}}$, i.e., where for any $\tilde{\mathbf{x}} \in \widehat{\text{ext}}(\mathcal{X}^{\text{LR}})$, the convex combination of $\tilde{\mathbf{x}}$ and $\hat{\mathbf{x}}$ must lie on the boundary of \mathcal{X}^{LR} . If $\tilde{\mathbf{x}} \in \widehat{\text{ext}}(\mathcal{X}^{\text{LR}})$, then we are done. Thus, we need to prove that when $\tilde{\mathbf{x}} \notin \widehat{\text{ext}}(\mathcal{X}^{\text{LR}})$, there does in fact exist $\bar{\mathbf{x}} \in \widehat{\text{ext}}(\mathcal{X}^{\text{LR}})$ for which $\mathbf{w}^\top \bar{\mathbf{x}} > \tau$.

We prove this by contradiction. Specifically, we assume that $\mathbf{w}^\top \mathbf{x} \leq \tau \forall \mathbf{x} \in \widehat{\text{ext}}(\mathcal{X}^{\text{LR}})$; we refer to this as *Assumption A1*. We begin by outlining the steps that lead to a contradiction, after which we prove the key step in detail:

1. Since $\tilde{\mathbf{x}} \in \mathcal{X}^{\text{LR}}$ and $\mathbf{w}^\top \tilde{\mathbf{x}} > \tau$, it must be true that $\tilde{\mathbf{x}} \in \mathcal{X}^{\text{LR}} \cap \{\mathbf{x} \mid \mathbf{w}^\top \mathbf{x} \geq \tau\}$.
2. We will show that this implies that $\tilde{\mathbf{x}} \in \text{conv}(\widehat{\text{ext}}(\mathcal{X}^{\text{LR}}) \cup \{\hat{\mathbf{x}}\}) \cap \{\mathbf{x} \mid \mathbf{w}^\top \mathbf{x} \geq \tau\}$.
3. This implies that $\tilde{\mathbf{x}} \in \text{conv}(\widehat{\text{ext}}(\mathcal{X}^{\text{LR}}) \cup \{\hat{\mathbf{x}}\})$.
4. Since $\tilde{\mathbf{x}} \notin \widehat{\text{ext}}(\mathcal{X}^{\text{LR}})$ and $\tilde{\mathbf{x}} \neq \hat{\mathbf{x}}$, this implies that $\tilde{\mathbf{x}}$ can be written strictly as a convex combination of other solutions in \mathcal{X}^{LR} , which contradicts the condition that $\tilde{\mathbf{x}}$ is an extreme point of \mathcal{X}^{LR} .

The key step in this proof is Step 2. To show this, we prove that when $\mathbf{w}^\top \mathbf{x} \leq \tau \forall \mathbf{x} \in \widehat{\text{ext}}(\mathcal{X}^{\text{LR}})$ (i.e., Assumption A1), it must be true that

$$\text{conv}(\widehat{\text{ext}}(\mathcal{X}^{\text{LR}}) \cup \{\hat{\mathbf{x}}\}) \cap \{\mathbf{x} \mid \mathbf{w}^\top \mathbf{x} \geq \tau\} = \mathcal{X}^{\text{LR}} \cap \{\mathbf{x} \mid \mathbf{w}^\top \mathbf{x} \geq \tau\}.$$

Note that the subset direction (\subseteq) must be true, as $\text{conv}(\widehat{\text{ext}}(\mathcal{X}^{\text{LR}}) \cup \{\hat{\mathbf{x}}\}) \subseteq \mathcal{X}^{\text{LR}}$. Thus, what remains is to prove the superset direction (\supseteq). For simplicity, let $\mathcal{G} := \{\mathbf{x} \mid \mathbf{w}^\top \mathbf{x} \geq \tau\}$.

To prove this, we will show that

$$\text{conv}(\widehat{\text{ext}}(\mathcal{X}^{\text{LR}}) \cup \{\hat{\mathbf{x}}\}) \cap \mathcal{G} = \widehat{\text{cone}}(\mathcal{X}^{\text{LR}}) \cap \mathcal{G}, \quad (\text{A4})$$

where

$$\widehat{\text{cone}}(\mathcal{X}^{\text{LR}}) = \{\hat{\mathbf{x}} + \sum_{\mathbf{x} \in \widehat{\text{ext}}(\mathcal{X}^{\text{LR}})} \lambda_{\mathbf{x}}(\mathbf{x} - \hat{\mathbf{x}}) \mid \lambda \geq \mathbf{0}\}$$

defines the cone generated by the extreme rays $\{\mathbf{y}(\lambda, \mathbf{x}, \hat{\mathbf{x}})\}_{\lambda \geq 0} \forall \mathbf{x} \in \widehat{\text{ext}}(\mathcal{X}^{\text{LR}})$, where $\mathbf{y}(\lambda, \mathbf{x}, \hat{\mathbf{x}}) := \hat{\mathbf{x}} + \lambda(\mathbf{y} - \hat{\mathbf{x}})$. Once we prove (A4), Step 2 must hold as $\mathcal{X}^{\text{LR}} \subseteq \widehat{\text{cone}}(\mathcal{X}^{\text{LR}})$ (Bodur et al. 2022).

We prove this by construction. For each $\mathbf{x}_i \in \widehat{\text{ext}}(\mathcal{X}^{\text{LR}})$, let \mathbf{t}^i denote the intersections between the extreme ray $\{\mathbf{y}(\lambda, \mathbf{x}_i, \hat{\mathbf{x}})\}_{\lambda \geq 0}$ and $\mathbf{w}^\top \mathbf{x} = \tau$. Note that by definition, $\widehat{\text{cone}}(\mathcal{X}^{\text{LR}}) \cap \mathcal{G} = \text{conv}(\mathbf{t}^1, \dots, \mathbf{t}^m, \hat{\mathbf{x}})$. From Assumption A1, $\mathbf{w}^\top \mathbf{x}^i < \tau$ and $\mathbf{w}^\top \hat{\mathbf{x}} > \tau$, which implies that this intersection must have happened on the line segment $\mathcal{L}(\mathbf{x}^i, \hat{\mathbf{x}}) := \{\mathbf{x} \mid \rho \mathbf{x}^i + (1 - \rho)\hat{\mathbf{x}}, \rho \in [0, 1]\}$, each of which is feasible in $\text{conv}(\widehat{\text{ext}}(\mathcal{X}^{\text{LR}}) \cup \{\hat{\mathbf{x}}\})$. Therefore, each point $\mathbf{t}^1, \dots, \mathbf{t}^m$ must also be in $\text{conv}(\widehat{\text{ext}}(\mathcal{X}^{\text{LR}}) \cup \{\hat{\mathbf{x}}\}) \cap \mathcal{G}$, which implies that $\widehat{\text{cone}}(\mathcal{X}^{\text{LR}}) \cap \mathcal{G} \subseteq \text{conv}(\widehat{\text{ext}}(\mathcal{X}^{\text{LR}}) \cup \{\hat{\mathbf{x}}\}) \cap \mathcal{G}$. Finally, we also know that $\text{conv}(\widehat{\text{ext}}(\mathcal{X}^{\text{LR}}) \cup \{\hat{\mathbf{x}}\}) \subseteq \mathcal{X}^{\text{LR}} \subseteq \widehat{\text{cone}}(\mathcal{X}^{\text{LR}})$. Thus, we have proved (A4), and Step 2; we are done with Part 1.

Proof of Part 2: Let $\bar{\mathbf{x}} \in \widehat{\text{ext}}(\mathcal{X}^{\text{LR}})$ be the adjacent extreme point to $\hat{\mathbf{x}}$ in \mathcal{X}^{LR} that satisfies $\mathbf{w}^\top \bar{\mathbf{x}} > \tau$, as proved in Part 1. By definition, $\hat{\mathbf{x}}$ and $\bar{\mathbf{x}}$ are defined by constraints that differ by only one active constraint, which must imply that there exists linearly independent $\mathbf{a}_1, \dots, \mathbf{a}_{n-1}$ in \mathbf{A} for which $\mathbf{a}_i^\top \hat{\mathbf{x}} = \mathbf{a}_i^\top \bar{\mathbf{x}} = b_i$ for all $i \in \{1, \dots, n-1\}$; recall we assumed that \mathcal{X}^{LR} is bounded (Bertsimas and Tsitsiklis 1997). Let $\hat{\boldsymbol{\theta}} = \sum_{i=1}^{n-1} \lambda_i \mathbf{a}_i$ where $\lambda_i \in (0, 1)$ for all $i \in \{1, \dots, n-1\}$. By definition, $\hat{\boldsymbol{\theta}}^\top \hat{\mathbf{x}} = \hat{\boldsymbol{\theta}}^\top \bar{\mathbf{x}}$ must hold and $\bar{\mathbf{x}}$ is the only solution in $\text{ext}(\mathcal{X}^{\text{LR}}) \setminus \{\hat{\mathbf{x}}\}$ that satisfies the equality. Furthermore, since $\hat{\boldsymbol{\theta}}^\top (\hat{\mathbf{x}} - \bar{\mathbf{x}}) = 0$, one of the constraints of $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}})$ is active, and $\hat{\boldsymbol{\theta}}$ must be on the boundary of $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}})$.

Proof of Part 3: We first show that there cannot exist any $\mathbf{x} \in \mathcal{X}^{\text{LR}}(\mathcal{W})$ for which $\hat{\boldsymbol{\theta}}^\top \hat{\mathbf{x}} = \hat{\boldsymbol{\theta}}^\top \mathbf{x}$. To do so, we define the line segment between $\hat{\mathbf{x}}$ and $\bar{\mathbf{x}}$ as $\mathcal{L} := \{\mathbf{y} \mid \mathbf{y} = \rho \hat{\mathbf{x}} + (1 - \rho)\bar{\mathbf{x}}, \rho \in [0, 1]\}$. Note that by the definition of $\hat{\boldsymbol{\theta}}$, $\mathcal{L} \subseteq \mathcal{X}^{\text{LR}}$ describes the complete set of solutions $\mathbf{x} \in \mathcal{X}^{\text{LR}}$ for which $\hat{\boldsymbol{\theta}}^\top \hat{\mathbf{x}} = \hat{\boldsymbol{\theta}}^\top \mathbf{x}$. In other words, $\hat{\boldsymbol{\theta}}^\top \mathbf{y} < \hat{\boldsymbol{\theta}}^\top \mathbf{x} \forall \mathbf{y} \in \mathcal{X}^{\text{LR}} \setminus \mathcal{L}$.

We now show that $\mathcal{L} \notin \mathcal{X}^{\text{LR}}(\mathcal{W})$. This is straightforward, as $\mathbf{w}^\top \bar{\mathbf{x}} > \tau$ and $\mathbf{w}^\top \hat{\mathbf{x}} > \tau$, which implies that $\mathbf{w}^\top \mathbf{y} > \tau$ for all $\mathbf{y} \in \mathcal{L}$. This implies that $\hat{\boldsymbol{\theta}}^\top (\hat{\mathbf{x}} - \mathbf{x}) > 0 \forall \mathbf{x} \in \mathcal{X}^{\text{LR}}(\mathcal{W})$, which by definition implies that $\hat{\boldsymbol{\theta}}$ is in the interior of $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}))$. Since $\boldsymbol{\theta}$ is on the boundary of $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}})$ but is not for $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}))$, it must be true that $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}) \subset \mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}))$. \square

A1.6. Proof of Proposition 3

When \mathbf{w} is on the boundary of $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}})$, it implies that there exists a $\bar{\mathbf{x}} \in \text{ext}(\mathcal{X}^{\text{LR}}) \setminus \{\hat{\mathbf{x}}\}$ such that $\mathbf{w}^\top \hat{\mathbf{x}} = \mathbf{w}^\top \bar{\mathbf{x}}$. Furthermore, if $\mathbf{w}^\top \mathbf{x} \leq \tau$ is a separating inequality, then $\mathbf{w}^\top \hat{\mathbf{x}} = \mathbf{w}^\top \bar{\mathbf{x}} > \tau$, which implies that $\bar{\mathbf{x}} \notin \mathcal{X}^{\text{LR}}(\mathcal{W})$. Theorem 3 implies that $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}) \subset \mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}))$. \square

A1.7. Proof of Theorem 4

We first provide a general overview of this proof, before giving the details. First, we note that $\mathcal{X}_1^{\text{LR}} := \text{conv}(\mathcal{X}^{\text{LR}}(\mathcal{W}_1) \cup \{\hat{\mathbf{x}}\})$ is also a valid linear relaxation for \mathcal{X} , since all of the extreme points of $\text{conv}(\mathcal{X})$ lie inside the set $\mathcal{X}_1^{\text{LR}}$. Next, note that $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}_1)) = \mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}_1^{\text{LR}})$. This is easy to see when we write out their definitions, i.e., note that

$$\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}_1)) := \{\boldsymbol{\theta} \mid \boldsymbol{\theta}^\top \hat{\mathbf{x}} \geq \boldsymbol{\theta}^\top \mathbf{x} \ \forall \mathbf{x} \in \mathcal{X}^{\text{LR}}(\mathcal{W}_1)\}$$

and

$$\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}_1^{\text{LR}}) := \{\boldsymbol{\theta} \mid \boldsymbol{\theta}^\top \hat{\mathbf{x}} \geq \lambda \boldsymbol{\theta}^\top \hat{\mathbf{x}} + (1 - \lambda) \boldsymbol{\theta}^\top \mathbf{x} \ \forall \mathbf{x} \in \mathcal{X}^{\text{LR}}(\mathcal{W}_1), \lambda \in [0, 1]\}.$$

Through simple algebra, we observe that the constraints in $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}_1^{\text{LR}})$ are equivalent to

$$(1 - \lambda) \boldsymbol{\theta}^\top \hat{\mathbf{x}} \geq (1 - \lambda) \boldsymbol{\theta}^\top \mathbf{x} \ \forall \mathbf{x} \in \mathcal{X}^{\text{LR}}(\mathcal{W}_1).$$

Since every inequality in $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}_1))$ appears in $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}_1^{\text{LR}})$, and every inequality in $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}_1^{\text{LR}})$ appears in $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}_1))$ (by multiplying both sides by a positive constant $\frac{1}{1-\lambda}$ when $\lambda \in [0, 1)$), the two sets must be equal.

Using these observations, we provide the proof of the theorem in three steps:

- Step 1: By definition, there must be an extreme point $\mathbf{x}^e \in \text{ext}(\mathcal{X}^{\text{LR}}) \setminus \{\hat{\mathbf{x}}\}$ where $\mathbf{x}^e \notin \mathcal{X}^{\text{LR}}(\mathcal{W}_1)$. We will show that there exists a $\boldsymbol{\theta}$ on the boundary of $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}_1))$ such that $\boldsymbol{\theta}^\top \mathbf{x}^e > \boldsymbol{\theta}^\top \hat{\mathbf{x}}$.
- Step 2: By definition, for this $\boldsymbol{\theta}$ there exists an extreme point $\mathbf{x}' \in \text{ext}(\mathcal{X}_1^{\text{LR}}) \setminus \{\hat{\mathbf{x}}\}$ where $\boldsymbol{\theta}^\top \mathbf{x}' = \boldsymbol{\theta}^\top \hat{\mathbf{x}}$. We will show that \mathbf{x}' must lie on the plane $\mathbf{w}^\top \mathbf{x} = \tau_1$, i.e., $\mathbf{w}^\top \mathbf{x}' = \tau_1$.
- Step 3: Since $\mathbf{w}^\top \mathbf{x}' = \tau_1$, this implies that $\mathbf{w}^\top \mathbf{x}' > \tau_2$ and that $\mathbf{x}' \notin \mathcal{X}^{\text{LR}}(\mathcal{W}_2)$. Furthermore, since $\mathcal{X}_1^{\text{LR}}$ defines a valid linear relaxation for \mathcal{X} , and since $\mathbf{x}' \in \text{ext}(\mathcal{X}_1^{\text{LR}})$, we can apply Theorem 3 to $\mathcal{X}_1^{\text{LR}}$. This implies that $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}_1)) \subset \mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}_2))$.

We now provide the detailed proofs of each step leading to the final result.

Proof of Step 1: Consider the cone pointed at $\hat{\mathbf{x}}$:

$$\text{cone}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}_1)) := \left\{ \hat{\mathbf{x}} + \sum_{\mathbf{x} \in \text{ext}(\mathcal{X}^{\text{LR}}(\mathcal{W}_1))} \lambda_{\mathbf{x}} (\mathbf{x} - \hat{\mathbf{x}}) \mid \lambda \geq \mathbf{0} \right\}. \quad (\text{A5})$$

Let $\mathbf{x}^e \in \text{ext}(\mathcal{X}^{\text{LR}}) \setminus \{\hat{\mathbf{x}}\}$ be an extreme point where $\mathbf{w}^\top \mathbf{x}^e > \tau_1$, which implies that $\mathbf{x}^e \notin \mathcal{X}^{\text{LR}}(\mathcal{W}_1)$. Furthermore, since the entire line segment $\mathcal{L}(\hat{\mathbf{x}}, \mathbf{x}^e) := \{\rho \hat{\mathbf{x}} + (1 - \rho) \mathbf{x}^e \mid \rho \in [0, 1]\}$ is not in $\mathcal{X}^{\text{LR}}(\mathcal{W}_1)$, i.e., $\mathcal{L}(\hat{\mathbf{x}}, \mathbf{x}^e) \cap \mathcal{X}^{\text{LR}}(\mathcal{W}_1) = \emptyset$, then $\mathbf{x}^e \notin \text{cone}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}_1))$. Let $\mathbf{D}(\mathbf{x} - \hat{\mathbf{x}}) \leq \mathbf{0}$ denote the facets that define $\text{cone}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}_1))$, and note that each row $\mathbf{d}_i \in \mathbf{D}$ must be a boundary point of $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}_1))$; namely, $\mathbf{d}_i^\top \hat{\mathbf{x}} \geq \mathbf{d}_i^\top \mathbf{x} \ \forall \mathbf{x} \in \mathcal{X}^{\text{LR}}(\mathcal{W}_1)$, and there exists $\mathbf{x}' \in \mathcal{X}^{\text{LR}}(\mathcal{W}_1)$ such that $\mathbf{d}_i^\top \hat{\mathbf{x}} = \mathbf{d}_i^\top \mathbf{x}'$. By definition, there must exist a facet $\mathbf{d}_i^\top \mathbf{x} \leq \mathbf{d}_i^\top \hat{\mathbf{x}}$ such that $\mathbf{d}_i^\top \mathbf{x}^e > \mathbf{d}_i^\top \hat{\mathbf{x}}$; otherwise

$\mathbf{x}^e \in \text{cone}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}_1))$. For this facet $\mathbf{d}_i^\top \mathbf{x} \leq \mathbf{d}_i^\top \hat{\mathbf{x}}$, the vector \mathbf{d}_i is an example of a $\boldsymbol{\theta}$ which is on the boundary of $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}_1))$, and for which $\boldsymbol{\theta}^\top \mathbf{x}^e > \boldsymbol{\theta}^\top \hat{\mathbf{x}}$.

Proof of Step 2: From Step 1, let $\mathbf{x}' \in \text{ext}(\mathcal{X}_1^{\text{LR}}) \setminus \{\hat{\mathbf{x}}\}$ be any extreme point of $\mathcal{X}_1^{\text{LR}}$ where $\boldsymbol{\theta}^\top \mathbf{x}' = \boldsymbol{\theta}^\top \hat{\mathbf{x}}$. We consider a proof by contradiction. Namely, suppose that $\mathbf{w}^\top \mathbf{x}' < \tau_1$. Let $\mathbf{x}^e \in \text{ext}(\mathcal{X}^{\text{LR}}) \setminus \{\hat{\mathbf{x}}\}$ denote the extreme point that is cut off by the separating inequality $\mathbf{w}^\top \mathbf{x} \leq \tau_1$, i.e., $\mathbf{w}^\top \mathbf{x}^e > \tau_1$. Since $\mathbf{x}^e, \mathbf{x}' \in \mathcal{X}^{\text{LR}}$, this implies that the line segment $\mathcal{L}(\mathbf{x}^e, \mathbf{x}') := \{\rho \mathbf{x}^e + (1 - \rho) \mathbf{x}' \mid \rho \in [0, 1]\}$ lies within \mathcal{X}^{LR} , and more importantly, that there exists a $\rho^* \in (0, 1)$ for which $\mathbf{w}^\top (\rho^* \mathbf{x}^e + (1 - \rho^*) \mathbf{x}') = \tau_1$. This implies that $(\rho^* \mathbf{x}^e + (1 - \rho^*) \mathbf{x}') \in \mathcal{X}^{\text{LR}}(\mathcal{W}_1)$. However, note that

$$\boldsymbol{\theta}^\top (\rho^* \mathbf{x}^e + (1 - \rho^*) \mathbf{x}') > \boldsymbol{\theta}^\top (\rho^* \hat{\mathbf{x}} + (1 - \rho^*) \mathbf{x}') = \boldsymbol{\theta}^\top \hat{\mathbf{x}},$$

which implies that $\boldsymbol{\theta}$ cannot be inverse-feasible. However, since $\boldsymbol{\theta} \in \mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}_1)) \subset \mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$, this must be a contradiction. Thus, \mathbf{x}' must satisfy $\mathbf{w}^\top \mathbf{x}' = \tau_1$.

Proof of Step 3: The proof of Step 3 is straightforward, and follows exactly the procedure outlined in the summary above. Specifically, we can treat $\mathcal{X}_1^{\text{LR}}$ as the existing linear relaxation of \mathcal{X} , and we know $\mathcal{X}^{\text{LR}}(\mathcal{W}_2)$ as the RLR generated by adding the separating inequality $\mathbf{w}^\top \mathbf{x} \leq \tau_2$ to $\mathcal{X}_1^{\text{LR}}$. Since there exists a $\mathbf{x}' \in \text{ext}(\mathcal{X}_1^{\text{LR}}) \setminus \{\hat{\mathbf{x}}\}$ where $\mathbf{w}^\top \mathbf{x}' > \tau_2$, this implies that $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}_1^{\text{LR}}) = \mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}_1)) \subset \mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}_2))$. This concludes the proof. \square

A1.8. Proof of Theorem 5

The proof of this theorem follows closely from the proofs of Theorems 3 and 4. First, recall that $\mathcal{X}_1^{\text{LR}} := \text{conv}(\mathcal{X}^{\text{LR}}(\mathcal{W}_1) \cup \{\hat{\mathbf{x}}\})$ defines a valid linear relaxation of \mathcal{X} , and, furthermore, $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}_1^{\text{LR}}) = \mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}_1))$; this was shown in the proof of Theorem 4. Now, since $\mathcal{X}_1^{\text{LR}}$ is a valid linear relaxation, we can simply replace \mathcal{X}^{LR} in Theorem 3 with $\mathcal{X}_1^{\text{LR}}$, and we obtain precisely the result outlined in Theorem 5. \square

A1.9. Proof of Proposition 4

Recall that at each outer iteration of Algorithm 1, the master problem (IO-RLR) solves the problem

$$\min_{\boldsymbol{\theta}, \mathbf{s}} f(\mathbf{s}) \tag{A6a}$$

$$\text{s.t. } \boldsymbol{\theta} \in \mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}^k)) \tag{A6b}$$

$$\boldsymbol{\theta} \in \mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X}) \tag{A6c}$$

$$\boldsymbol{\theta} \in \boldsymbol{\Theta}(\mathbf{s}) \tag{A6d}$$

$$\mathbf{s} \in \mathcal{S}. \tag{A6e}$$

Let $(\boldsymbol{\theta}^k, \mathbf{s}^k)$ denote the optimal solution to this problem. Since $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}^k)) \subseteq \mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}^{k-1}))$ $\forall k \geq 0$, it must be true that $f(\mathbf{s}^k) \leq f(\mathbf{s}^{k-1})$. It thus suffices that show that if $f(\mathbf{s}^k) = f(\mathbf{s}^{k-1})$, then $(\boldsymbol{\theta}^k, \mathbf{s}^k)$ must be an optimal solution to (IO).

To prove the above statement, we consider a proof by contradiction. Suppose that $f(\mathbf{s}^k) = f(\mathbf{s}^{k-1})$, but that $f(\mathbf{s}^{k-1}) > f(\mathbf{s}^*)$ where \mathbf{s}^* is the optimal solution to the inverse problem. Recall that $\Theta(\mathbf{s})$ defines a linear mapping from \mathbf{s} to $\boldsymbol{\theta}$; we will write this as $\boldsymbol{\theta} = \mathbf{H}\mathbf{s} + \mathbf{d}$ where $\mathbf{s} \in \mathbb{R}^m$. We prove that this is a contradiction in two steps.

Step 1: We first remark that $\boldsymbol{\theta}^{k-1} := \mathbf{H}\mathbf{s}^{k-1} + \mathbf{d}$ must be in the interior of $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}^k))$. Specifically, $\mathcal{X}^{\text{LR}}(\mathcal{W}^k)$ must include a restricting inequality $(\boldsymbol{\theta}^{k-1})^\top \mathbf{x} \leq \tau$, where $\tau := \boldsymbol{\theta}^\top \hat{\mathbf{x}} - \eta$ for some depth parameter $\eta > 0$. This implies that $(\boldsymbol{\theta}^{k-1})^\top \hat{\mathbf{x}} > (\boldsymbol{\theta}^{k-1})^\top \mathbf{x} \quad \forall \mathbf{x} \in \mathcal{X}(\mathcal{W}_k)$, which, by definition, implies that $\boldsymbol{\theta}^{k-1}$ must be in the interior of $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}^k))$. This last observation is meaningful because it implies that for any bounded vector $\hat{\mathbf{s}} \in \mathbb{R}^m$, there exists $\hat{\lambda} > 0$ such that the vector $\boldsymbol{\theta}^{k-1} + \lambda(\mathbf{H}\hat{\mathbf{s}}) := \mathbf{H}(\mathbf{s}^{k-1} + \lambda\hat{\mathbf{s}}) + \mathbf{d}$ remains within $\mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}^k))$ for all $\lambda \in [0, \hat{\lambda}]$. We will use this fact at the end of Step 2.

Step 2: We now show that if $f(\mathbf{s}^{k-1}) > f(\mathbf{s}^*)$, then there must exist some $\lambda \in (0, 1]$ for which $\tilde{\mathbf{s}} := (1 - \lambda)\mathbf{s}^{k-1} + \lambda\mathbf{s}^*$ is also feasible to the problem (A6). Note that since $f(\mathbf{s})$ is a convex function,

$$f(\tilde{\mathbf{s}}) = f((1 - \lambda)\mathbf{s}^{k-1} + \lambda\mathbf{s}^*) \leq (1 - \lambda)f(\mathbf{s}^{k-1}) + \lambda f(\mathbf{s}^*) < f(\mathbf{s}^{k-1}).$$

Thus, if we can show the existence of such a feasible $\tilde{\mathbf{s}}$, then \mathbf{s}^k cannot be the optimal solution to (A6), which is a contradiction.

We now show the existence of such a feasible $\tilde{\mathbf{s}}$. First, we note that if both \mathbf{s}^{k-1} and \mathbf{s}^* map to inverse-feasible vectors, then it must be true that its convex combination also maps to inverse-feasible vectors, i.e., $\mathbf{H}(\rho\mathbf{s}^* + (1 - \rho)\mathbf{s}^{k-1}) + \mathbf{d} = \rho\boldsymbol{\theta}^* + (1 - \rho)\boldsymbol{\theta}^{k-1} \in \mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X}) \quad \forall \rho \in [0, 1]$; this is true because $\mathcal{R}_{\text{inv}}(\hat{\mathbf{x}}, \mathcal{X})$ is a cone. This observation implies that any convex combination of \mathbf{s}^{k-1} and \mathbf{s}^* must satisfy constraint (A6c). Next, since \mathcal{S} is convex, any convex combination of $\mathbf{s}^{k-1} \in \mathcal{S}$ and $\mathbf{s}^* \in \mathcal{S}$ must also satisfy (A6e). The only remaining constraint to discuss is (A6b). We now use the result from Step 1, where, for $\hat{\mathbf{s}} := \mathbf{s}^* - \mathbf{s}^{k-1}$, that there must be a feasible $\lambda \in (0, 1]$ such that $\mathbf{H}(\mathbf{s}^{k-1} + \lambda\hat{\mathbf{s}}) + \mathbf{d} \in \mathcal{K}(\hat{\mathbf{x}}, \mathcal{X}^{\text{LR}}(\mathcal{W}^k))$. If we define $\tilde{\mathbf{s}} := \mathbf{s}^{k-1} + \lambda\hat{\mathbf{s}} = (1 - \lambda)\mathbf{s}^{k-1} + \lambda\mathbf{s}^*$, then we have found a $\tilde{\mathbf{s}}$ that is feasible to (A6) and where $f(\tilde{\mathbf{s}}) < f(\mathbf{s}^k)$. Thus, \mathbf{s}^k cannot be optimal to (A6), which is a contradiction. Algorithm 1 must therefore terminate at the optimal solution. \square

A2. Additional Material for Section 4

A2.1. Computing separating inequalities for integer programs

Proposition A1. *Given $\mathbf{w} \in \mathbb{R}^n$ and $\hat{\mathbf{x}} \in \mathcal{X}$ where $\mathcal{X} \subset \mathbb{Z}^n$ is bounded, the problem*

$$\max_{\mathbf{x}} \{ \mathbf{w}^\top \mathbf{x} \mid \mathbf{x} \in \mathcal{X} \setminus \{\hat{\mathbf{x}}\} \}$$

can be solved as

$$\max_{\mathbf{x}, \mathbf{z}, \boldsymbol{\delta}} \mathbf{w}^\top \mathbf{x} \tag{A7a}$$

$$\text{s.t. } \mathbf{x} \in \mathcal{X} \tag{A7b}$$

$$z_i \leq x_i - \hat{x}_i + M\delta_i \quad \forall i \in \{1, \dots, n\} \tag{A7c}$$

$$z_i \leq \hat{x}_i - x_i + M(1 - \delta_i) \quad \forall i \in \{1, \dots, n\} \tag{A7d}$$

$$\sum_{i=1}^n z_i \geq 1 \tag{A7e}$$

$$\mathbf{z} \geq \mathbf{0}, \boldsymbol{\delta} \in \{0, 1\}^n \tag{A7f}$$

where M denotes a large constant for which $|\mathbf{x}^1 - \mathbf{x}^2| \leq M$, $\forall \mathbf{x}^1, \mathbf{x}^2 \in \mathcal{X}$.

Proof of Proposition A1. When $\mathcal{X} \subseteq \mathbb{Z}^n$ is bounded, the set $\mathcal{X} \setminus \{\hat{\mathbf{x}}\}$ can be written as

$$\mathcal{S}_A := \left\{ \mathbf{x} \in \mathcal{X} \mid \sum_{i=1}^n |x_i - \hat{x}_i| \geq 1 \right\}, \tag{A8a}$$

since any $\mathbf{x} \in \mathcal{X} \setminus \{\hat{\mathbf{x}}\}$ must satisfy $|\mathbf{x} - \hat{\mathbf{x}}|_1 \geq 1$. \mathcal{S}_A is also equivalent to

$$\mathcal{S}_B := \left\{ \mathbf{x} \in \mathcal{X} \mid \sum_{i=1}^n z_i \geq 1, 0 \leq z_i \leq |x_i - \hat{x}_i| \quad \forall i \in \{1, \dots, n\} \right\}. \tag{A9a}$$

It is easy to verify that for any $\mathbf{x} \in \mathcal{S}_A$, there exists a \mathbf{z} for which $(\mathbf{x}, \mathbf{z}) \in \mathcal{S}_B$, and similarly, for any $(\mathbf{x}, \mathbf{z}) \in \mathcal{S}_B$, the value of \mathbf{x} must satisfy $\mathbf{x} \in \mathcal{S}_A$. Furthermore, \mathcal{S}_B is by definition equivalent to

$$\mathcal{S}_C := \left\{ \mathbf{x} \in \mathcal{X} \mid \sum_{i=1}^n z_i \geq 1, 0 \leq z_i \leq \max\{x_i - \hat{x}_i, \hat{x}_i - x_i\} \quad \forall i \in \{1, \dots, n\} \right\}. \tag{A10a}$$

In \mathcal{S}_C , the constraint

$$0 \leq z_i \leq \max\{x_i - \hat{x}_i, \hat{x}_i - x_i\} \quad \forall i \in \{1, \dots, n\}$$

can be reformulated as

$$\begin{aligned} z_i &\leq x_i - \hat{x}_i + M\delta_i \quad \forall i \in \{1, \dots, n\} \\ z_i &\leq \hat{x}_i - x_i + M(1 - \delta_i) \quad \forall i \in \{1, \dots, n\} \\ z_i &\geq 0 \quad \forall i \in \{1, \dots, n\} \\ \delta_i &\in \{0, 1\} \quad \forall i \in \{1, \dots, n\}. \end{aligned}$$

We note that this is a valid reformulation by considering all three cases: (i) when $x_i - \hat{x}_i = 0$, $z_i = 0$; (ii) when $x_i - \hat{x}_i \geq 1$, $\delta_i = 0$ and $z_i \leq x_i - \hat{x}_i$; (iii) when $\hat{x}_i - x_i \geq 1$, $\delta_i = 1$ and $z_i \leq \hat{x}_i - x_i$. \square

A2.2. Details of Example 3

This section provides additional details on Example 3. To generate the approximations of the inverse-feasible region, we consider $\mathcal{K}(\hat{\mathbf{x}}, \tilde{\mathcal{X}})$ where $\tilde{\mathcal{X}}$ is defined by the standard linear relaxation (LR), a tighter linear relaxation (LR⁺), and the four different RLRs detailed in Example 3:

- **LR** (\mathcal{X}^{LR}): $\{\mathbf{x} \mid \mathbf{A}\mathbf{x} \leq \mathbf{b}, \mathbf{x} \in [0, 1]^n\}$
- **LR⁺**: $\{\mathbf{x} \in \mathcal{X}^{\text{LR}} \mid \hat{\boldsymbol{\theta}}^\top \mathbf{x} \leq \hat{\boldsymbol{\theta}}^\top \hat{\mathbf{x}}\}$
- **RLR0.5**: $\{\mathbf{x} \in \mathcal{X}^{\text{LR}} \mid \hat{\boldsymbol{\theta}}^\top \mathbf{x} \leq \hat{\boldsymbol{\theta}}^\top \hat{\mathbf{x}} - 0.5\eta^*\}$
- **RLR1**: $\{\mathbf{x} \in \mathcal{X}^{\text{LR}} \mid \hat{\boldsymbol{\theta}}^\top \mathbf{x} \leq \hat{\boldsymbol{\theta}}^\top \hat{\mathbf{x}} - \eta^*\}$
- **RLR2**: $\{\mathbf{x} \in \mathcal{X}^{\text{LR}} \mid \hat{\boldsymbol{\theta}}^\top \mathbf{x} \leq \hat{\boldsymbol{\theta}}^\top \hat{\mathbf{x}} - 2\eta^*\}$
- **RLR3**: $\{\mathbf{x} \in \mathcal{X}^{\text{LR}} \mid \hat{\boldsymbol{\theta}}^\top \mathbf{x} \leq \hat{\boldsymbol{\theta}}^\top \hat{\mathbf{x}} - 3\eta^*\}$

As mentioned in the example, the value $\eta^* := \hat{\boldsymbol{\theta}}^\top \hat{\mathbf{x}} - \tau^*(\hat{\boldsymbol{\theta}})$ in Figure 4 denotes the depth of the tight separating inequality generated by solving Problem (14) with $\mathbf{w} = \hat{\boldsymbol{\theta}}$. By definition, **RLR0.5** and **RLR1** define pure-RLRs, whereas **RLR2** and **RLR3** define mixed-RLRs.

Table A1 provides the sensitivity ranges of each objective coefficient θ_i . The leftmost column provides the nominal values of $\hat{\boldsymbol{\theta}}$ and $\hat{\mathbf{x}}$, whereas the rightmost column provides the true sensitivity ranges. The remaining columns provide the ranges generated by the linear relaxations and RLRs.

Table A1: Sensitivity ranges for all θ_i for different linear relaxations

index	$\hat{\theta}_i$	\hat{x}_i	LR		LR ⁺		RLR0.5		RLR1		RLR2		RLR3		TRUE	
			Δ_i^-	Δ_i^+	Δ_i^-	Δ_i^+	Δ_i^-	Δ_i^+	Δ_i^-	Δ_i^+	Δ_i^-	Δ_i^+	Δ_i^-	Δ_i^+	Δ_i^-	Δ_i^+
0	560	0	-	-	$-\infty$	0	$-\infty$	6.50	$-\infty$	13	$-\infty$	26	$-\infty$	39	$-\infty$	43
1	1125	0	-	-	$-\infty$	0	$-\infty$	7.11	$-\infty$	14	$-\infty$	27.17	$-\infty$	39.59	$-\infty$	189
2	300	0	-	-	$-\infty$	0	$-\infty$	7.26	$-\infty$	13.53	$-\infty$	26	$-\infty$	39	$-\infty$	62
3	620	1	-	-	0	∞	-6.50	∞	-13	∞	-26	∞	-39	∞	-74	∞
4	2100	0	-	-	$-\infty$	0	$-\infty$	28.57	$-\infty$	53.22	$-\infty$	93.77	$-\infty$	125.68	$-\infty$	809
5	431	1	-	-	0	∞	-11.72	∞	-23.44	∞	-41	∞	-54.66	∞	-140	∞
6	68	0	-	-	$-\infty$	0	$-\infty$	6.50	$-\infty$	13	$-\infty$	26	$-\infty$	39	$-\infty$	19
7	328	1	-	-	0	∞	-6.50	∞	-13	∞	-26	∞	-39	∞	-85	∞
8	47	1	-	-	0	∞	-6.50	∞	-13	∞	-26	∞	-39	∞	-27	∞
9	122	0	-	-	$-\infty$	0	$-\infty$	6.50	$-\infty$	13	$-\infty$	26	$-\infty$	39	$-\infty$	33
10	322	1	-	-	0	∞	-6.50	∞	-13	∞	-26	∞	-39	∞	-87	∞
11	196	1	-	-	0	∞	-6.50	∞	-13	∞	-26	∞	-39	∞	-19	∞
12	41	1	-	-	0	∞	-6.50	∞	-13	∞	-26	∞	-39	∞	-32	∞
13	25	0	-	-	$-\infty$	0	$-\infty$	6.50	$-\infty$	13	$-\infty$	26	$-\infty$	39	$-\infty$	13
14	425	1	-	-	0	∞	-18.87	∞	-35.86	∞	-62.57	∞	-83.27	∞	-248	∞
15	4260	1	-	-	0	∞	-26.63	∞	-53.26	∞	-98.37	∞	-137.19	∞	-1389	∞
16	416	1	-	-	0	∞	-19.45	∞	-38.90	∞	-67.87	∞	-90.32	∞	-245	∞

Continued on next page

Table A1 – continued from previous page

index	$\hat{\theta}_i$	\hat{x}_i	LR		LR ⁺		RLR0.5		RLR		RLR (-2η)		RLR (-3η)		TRUE	
			Δ_i^-	Δ_i^+	Δ_i^-	Δ_i^+	Δ_i^-	Δ_i^+	Δ_i^-	Δ_i^+	Δ_i^-	Δ_i^+	Δ_i^-	Δ_i^+	Δ_i^-	Δ_i^+
17	115	0	-	-	$-\infty$	0	$-\infty$	6.50	$-\infty$	13	$-\infty$	26	$-\infty$	39	$-\infty$	41
18	82	1	-	-	0	∞	-6.50	∞	-13	∞	-26	∞	-39	∞	-44	∞
19	22	1	-	-	0	∞	-6.50	∞	-13	∞	-26	∞	-39	∞	-13	∞
20	631	0	-	-	$-\infty$	0	$-\infty$	6.50	$-\infty$	13	$-\infty$	26	$-\infty$	39	$-\infty$	72
21	132	0	-	-	$-\infty$	0	$-\infty$	6.50	$-\infty$	13	$-\infty$	26	$-\infty$	39	$-\infty$	57
22	420	1	-	-	0	∞	-6.50	∞	-13	∞	-26	∞	-39	∞	-76	∞
23	86	0	-	-	$-\infty$	0	$-\infty$	6.50	$-\infty$	13	$-\infty$	26	$-\infty$	39	$-\infty$	39
24	42	1	-	-	0	∞	-6.50	∞	-13	∞	-26	∞	-39	∞	-13	∞
25	103	1	-	-	0	∞	-6.50	∞	-13	∞	-26	∞	-39	∞	-19	∞
26	215	1	-	-	0	∞	-6.50	∞	-13	∞	-26	∞	-39	∞	-75	∞
27	81	1	-	-	0	∞	-6.50	∞	-13	∞	-26	∞	-39	∞	-29	∞
28	91	1	-	-	0	∞	-6.50	∞	-13	∞	-26	∞	-39	∞	-71	∞
29	26	0	-	-	$-\infty$	0	$-\infty$	6.50	$-\infty$	13	$-\infty$	26	$-\infty$	39	$-\infty$	13
30	49	1	-	-	0	∞	-6.50	∞	-13	∞	-26	∞	-39	∞	-43	∞
31	420	1	-	-	0	∞	-18.36	∞	-36.72	∞	-64.07	∞	-85.23	∞	-285	∞
32	316	0	-	-	$-\infty$	0	$-\infty$	6.50	$-\infty$	13	$-\infty$	26	$-\infty$	39	$-\infty$	18
33	72	1	-	-	0	∞	-6.50	∞	-13	∞	-26	∞	-39	∞	-31	∞
34	71	1	-	-	0	∞	-6.50	∞	-13	∞	-26	∞	-39	∞	-13	∞
35	49	1	-	-	0	∞	-6.50	∞	-13	∞	-26	∞	-39	∞	-19	∞
36	108	1	-	-	0	∞	-6.50	∞	-13	∞	-26	∞	-39	∞	-58	∞
37	116	1	-	-	0	∞	-6.50	∞	-13	∞	-26	∞	-39	∞	-19	∞
38	90	1	-	-	0	∞	-6.50	∞	-13	∞	-26	∞	-39	∞	-64	∞
39	738	1	-	-	0	∞	-30.79	∞	-61.57	∞	-107.54	∞	-143.17	∞	-382	∞
40	1811	1	-	-	0	∞	-65	∞	-130	∞	-226.81	∞	-301.70	∞	-884	∞
41	430	1	-	-	0	∞	-19.39	∞	-38.77	∞	-67.63	∞	-89.96	∞	-286	∞
42	3060	1	-	-	0	∞	-148.41	∞	-296.83	∞	-518.42	∞	-690.16	∞	-2076	∞
43	215	1	-	-	0	∞	-6.50	∞	-13	∞	-26	∞	-39	∞	-38	∞
44	58	0	-	-	$-\infty$	0	$-\infty$	6.50	$-\infty$	13	$-\infty$	26	$-\infty$	39	$-\infty$	13
45	296	0	-	-	$-\infty$	0	$-\infty$	8.09	$-\infty$	16.18	$-\infty$	28.26	$-\infty$	39	$-\infty$	77
46	620	1	-	-	0	∞	-14.73	∞	-29.45	∞	-51.54	∞	-68.77	∞	-181	∞
47	418	1	-	-	0	∞	-10.83	∞	-21.65	∞	-37.89	∞	-50.53	∞	-190	∞
48	47	1	-	-	0	∞	-6.50	∞	-13	∞	-26	∞	-39	∞	-18	∞
49	81	1	-	-	0	∞	-6.50	∞	-13	∞	-26	∞	-39	∞	-61	∞

Note that the two linear relaxations **LR** and **LR**⁺ are entirely uninformative when used for sensitivity analysis. For **LR**, we are not able to generate any ranges on $\hat{\theta}$ simply because $\hat{\theta}$ does not render \hat{x} optimal over the relaxation of Problem (11). For the tighter linear relaxation **LR**⁺, $\hat{\theta}$ defines a normal vector of a facet tangent to \hat{x} , and any perturbations on $\hat{\theta}$ in a non-trivial

direction (i.e., negative direction when $x_i = 1$ or positive direction when $x_i = 0$) will not retain the optimality of \hat{x} over this relaxation. Figure A1 provides an illustration.

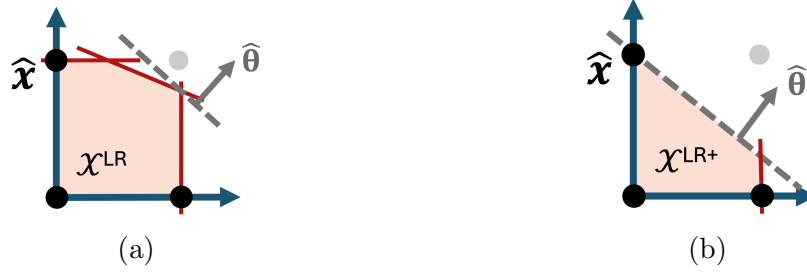


Figure A1 An illustration of why \mathcal{X}^{LR} and $\mathcal{X}^{\text{LR}+}$ may not lead to informative sensitivity ranges. Specifically, any perturbation of $\hat{\theta}$ in a non-trivial direction will no longer render \hat{x} optimal over either relaxations.

We now outline how we calculated the average coverage and level of overapproximation that is shown in Figure 4. The percentage of coverage of sensitivity ranges for each **RLRK** where $\mathbf{K} \in \{0.5, 1, 2, 3\}$ is calculated using the formula

$$\frac{\sum_{i \in n} (\Delta_i^{+(\text{TRUE})} - \Delta_i^{+(\text{RLRK})}) + (\Delta_i^{-(\text{TRUE})} - \Delta_i^{-(\text{RLRK})})}{\sum_{i \in n} (\Delta_i^{+(\text{TRUE})} + \Delta_i^{-(\text{TRUE})})}$$

where any value of ∞ or $-\infty$ is replaced with the value of 0, as one direction will always be unbounded when $x_i \in \{0, 1\}$ is a binary variable. Similarly, the percentages of overapproximation can be calculated using the formula

$$\frac{\sum_{i \in n} v_i^k}{\sum_{i \in n} (\Delta_i^{+(\text{TRUE})} + \Delta_i^{-(\text{TRUE})})}$$

where each v_i^k tabulates the overapproximation (if any) of **RLRK** for each coordinate $i \in n$, i.e.,

$$v_i^k = \begin{cases} \Delta_i^{+(\text{RLRK})} - \Delta_i^{+(\text{TRUE})} & \text{if } \Delta_i^{+(\text{RLRK})} - \Delta_i^{+(\text{TRUE})} > 0 \\ \Delta_i^{-(\text{RLRK})} - \Delta_i^{-(\text{TRUE})} & \text{if } \Delta_i^{-(\text{RLRK})} - \Delta_i^{-(\text{TRUE})} > 0 \\ 0 & \text{otherwise.} \end{cases}$$

A3. Additional Material for Section 5

A3.1. Details of the adaptive policy

With a poorly specified static depth selection policy, the RLR-CP Algorithm may either degenerate into the standard cutting plane algorithm when the depth of restricting inequalities are too large, or become overly conservative when the depths are too shallow. We propose an adaptive policy that allows for real-time modifications to the depth parameters chosen by the static policy. This can both mitigate the algorithm's sensitivity to the static policy and enhance its overall performance.

In its most general form, the adaptive policy adjusts the depth of the current or subsequent restricting inequality based on the time spent in the cutting plane subroutine. Specifically, as more time is spent within the subroutine, the depth of the most recently added restricting inequality can be progressively decreased. Conversely, if the subroutine is largely bypassed – indicating that the depth is too conservative – the next restricting inequality that is added can be initialized with a larger depth value. To implement this adaptive policy, we require introducing an additional t index to η^k , where η_t^k now denotes the depth of the k -th restricting inequality at iteration t in the cutting plane subroutine.

To address the first case, we can define an adaptive policy that decreases the depth parameter of the most recently added restricting inequality after each (or several) iteration(s) of this subroutine. One simple example would be to adjust η_t^k as

$$\eta_t^k = \frac{\eta_0^k}{t},$$

where the depth of the restricting inequality $(\theta^k)^\top \mathbf{x} \leq (\theta^k)^\top \hat{\mathbf{x}} - \eta_t^k$ is iteratively made more shallow as cuts are generated by the cutting plane subroutine.

To address the second case, we can define an adaptive policy that modifies the baseline depth of the next restricting inequality based on the number of iterations spent in the cutting plane subroutine under the current restricting inequality. Specifically, recall from equation (19) that a fixed $\tilde{\eta}$ is used to define a static policy where $\eta_{t=0}^{k+1} = \tilde{\eta} |(\theta^{k+1})^\top \hat{\mathbf{x}}|$. To implement an adaptive policy, we can simply replace $\tilde{\eta}$ with $\tilde{\eta}^{k+1}$ where $\tilde{\eta}^{k+1}$ is a decreasing function in the number of cuts added by the subroutine in iteration k , which we denote as T^k . For example, we can define $\tilde{\eta}^k$ as

$$\tilde{\eta}^{k+1} = \frac{\alpha}{T^k} \tilde{\eta}^k$$

for some value of $\alpha > 1$.

We conclude by noting that the main purpose of this section is to provide a conceptual framework for defining this adaptive policy, and there may be many alternative specifications for how we can adjust the depth of restricting inequalities based on real-time performance.

A4. Additional Material for Section 6.2

A4.1. Generation of forward problem instances

The multi-objective knapsack problems that we consider in the numerical experiments can be written as

$$\max_{\mathbf{x}} \left\{ \sum_{j=1}^{10} s_j \mathbf{f}_j \mathbf{x} \mid \mathbf{A} \mathbf{x} \leq \mathbf{b}, \mathbf{x} \in \{0, 1\}^n \right\}$$

where $n = \{100, 200, 400, 800\}$. For each n , the knapsack problem is defined using a constraint matrix $\mathbf{A} \in \mathbb{R}^{5 \times n}$ consisting of randomly generated integers between 0 and 201, a vector \mathbf{b} that is generated by taking 40% of the sum of entries of each row of \mathbf{A} , and vectors \mathbf{f}_i that are composed of randomly generated integers between 50 and 200.

A4.2. Details of CP algorithms

We consider two cutting plane algorithms as benchmarks, one of which serves as the cutting plane subroutine in our RLR-CP algorithm. Note that the main difference between these benchmarks and our RLR-CP algorithm is that the benchmarks simply do not contain constraint $\theta \in \mathcal{K}(\hat{x}, \mathcal{X}(\mathcal{W}^k))$. Specifically, the standard CP algorithm (Wang 2009) is defined precisely by the subroutine given in IO-RLR of Algorithm 1 when we remove the constraint $\theta \in \mathcal{K}(\hat{x}, \mathcal{X}(\mathcal{W}^k))$ from CP-M. Similarly, the trust-region based CP algorithm (Bodur et al. 2022) differs only in the definition of IO-RLR; instead of solving

$$\max\{(\tilde{\theta}^t)^\top x \mid x \in \mathcal{X}\}, \quad (\text{A11})$$

at each iteration, the trust-region based CP algorithm solves

$$\max\{(\tilde{\theta}^t)^\top x \mid x \in \mathcal{X} \cap \text{TR}(\hat{x}, \delta)\} \quad (\text{A12})$$

where $\text{TR}(\hat{x}, \delta) = \{x \mid \|x - \hat{x}\|_1 \leq \delta\}$ and δ controls the size of the L_1 -norm ball centered at \hat{x} .

We implemented the same version of the trust-region based CP algorithm that was implemented in Bodur et al. (2022): starting with $\delta^{(0)} = 1$ where $k = 0$, we solve the CP subroutine where δ in (A12) is replaced with $\delta^{(k)} = 2^k$. When cuts can no longer be generated using $\delta^{(k)}$ (i.e., there does not exist any $x^t \in \mathcal{X} \cap \text{TR}(\hat{x}, \delta^{(k)})$ such that $(\tilde{\theta}^t)^\top x^t > (\tilde{\theta}^t)^\top \hat{x}$), we increase the value of k to $k + 1$, doubling the value of δ . Then, for every 10 cuts that are generated, we force $\delta = \infty$ for one iteration, which enables the algorithm to terminate if it cannot find any violated cut in all of \mathcal{X} .

Finally, to speed up both the benchmark CP algorithms and the CP subroutines, we consider the early-stop heuristic from Bodur et al. (2022), where, rather than solving (A11) and (A12) to optimality at each iteration, we can instead stop the solver earlier when two conditions are both satisfied: (i) at least 5 seconds have elapsed and (ii) the best solution found satisfies the condition $(\theta^k)^\top x^t > (\theta^k)^\top \hat{x}$. This early-stop heuristic enables a more efficient cut generation process for both the CP and RLR-CP algorithms.

A4.3. Additional results and insights

In Section 6, we examined the aggregated results for the multi-objective knapsack problems with 200 and 400 variables. For completeness, we provide the results for the 100 and 800 variable problems in Tables A2 and A3. We note that the three key insights outlined in Section 6.2 remain consistent across these additional instances.

Next, we provide further analysis on the depth parameters used in the numerical experiments. Specifically, we compare the value of the best performing depth parameters to those that would have been generated if the algorithm had only computed tight separating inequalities. The purpose

Time Limit	5s				10s				20s			
Optimality Gap	15%	5%	1%	0%	15%	5%	1%	0%	15%	5%	1%	0%
RLR-CP: $\tilde{\eta} = 0.0001$	0	0	0	0	0	0	0	0	0	0	0	0
RLR-CP: $\tilde{\eta} = 0.001$	7	2	0	0	10	9	6	2	10	10	10	9
RLR-CP: $\tilde{\eta} = 0.01$	6	5	5	5	7	7	7	7	10	10	10	10
RLR-CP: $\tilde{\eta} = 0.1$	1	1	1	1	5	5	5	5	8	8	8	8
CP: Standard	0	0	0	0	4	4	4	4	7	7	7	7
CP: Trust Region	1	1	1	1	4	4	4	4	8	8	8	8

Table A2 Counts of 100 variable instances reaching specified optimality gaps within time limits.

Time Limit	500s				1500s				3000s			
Optimality Gap	15%	5%	1%	0%	15%	5%	1%	0%	15%	5%	1%	0%
RLR-CP: $\tilde{\eta} = 0.00001$	7	4	2	0	10	6	5	0	10	10	5	0
RLR-CP: $\tilde{\eta} = 0.0001$	6	3	1	0	7	6	3	2	8	7	5	2
RLR-CP: $\tilde{\eta} = 0.001$	0	0	0	0	6	6	3	0	9	9	6	2
RLR-CP: $\tilde{\eta} = 0.01$	0	0	0	0	3	3	2	0	7	7	6	2
RLR-CP: $\tilde{\eta} = 0.1$	0	0	0	0	3	3	1	0	8	8	3	1
CP: Standard	0	0	0	0	0	0	0	0	8	8	8	8
CP: Trust Region	0	0	0	0	2	2	2	2	8	8	8	8

Table A3 Counts of 800 variable instances reaching specified optimality gaps within time limits.

of this comparison is to understand whether the insights illustrated in Example 3 hold, i.e., whether generating separating inequalities may be overly conservative.

To derive these insights, we apply Algorithm 1 to the same IO instances by generating only separating inequalities. Specifically, given any candidate θ^k generated at an iteration k in this new implementation, we solve Problem (14) in RLR-UPDATE to generate a tight separating inequality $(\theta^k)^\top \mathbf{x} \leq \tau^*(\theta^k)$, stopping only when it is no longer possible to do so, i.e., when $(\theta^{\tilde{k}})^\top \hat{\mathbf{x}} = \tau^*(\theta^{\tilde{k}})$ for some iteration \tilde{k} .

	100 Vars.	200 Vars.	400 Vars.	800 Vars.
Avg. absolute depth (η^k)	2.25	2.02	1.43	0.2
Avg. relative depth ($\eta^k / (\theta^k)^\top \hat{\mathbf{x}} $)	0.00034	0.00016	0.00005	0.000004

Table A4 Average depth parameter values for $k < \tilde{k}$, averaged across 10 instances for model size.

Table A4 shows the depth parameter values, averaged across all iterations $k < \tilde{k}$ and all problem instances for the different variable sizes. The first row shows the average depth in absolute values, while the second row shows the average depth in terms of the multiplicative value $\tilde{\eta}^k$ for which $\tilde{\eta}^k |(\theta^k)^\top \hat{\mathbf{x}}|$ would have to lead to this same value. If we compare the values of the second row to

those corresponding to the best average-performing choice of $\tilde{\eta}$ in Tables A2 (100 variables), 2 (200 variables), 3 (400 variables), and A3 (800 variables), we find that the best performing choices of $\tilde{\eta}$ are larger than the values in the second row. To pick out one example, we can consider the 200 variable instances, shown in Table 2. In these instance, setting a depth parameter of $\tilde{\eta} = 0.0001$, which is almost identical to the corresponding value in the second row of Table A4, leads to an algorithm that performs much worse than if $\tilde{\eta} = 0.001$. This insight is consistent with what was observed in Example 3, i.e., that separating inequalities may be overly conservative and that deeper restricting inequalities – albeit not excessively deep – can cover significantly more of the inverse-feasible region and result in better performing algorithms. When we increase $\tilde{\eta}$ even further to $\tilde{\eta} = 0.01$, the performance of the algorithm degrades again. Similar insights can be observed across all instances. In summary, these observation reinforce the idea that the best performing specification of Algorithm 1 are those that strike a balance between using RLRs and the cutting plane subroutine, rather than trying to avoid the latter altogether.