informs
https://pubsonline.informs.org/journal/mnsc

MANAGEMENT SCIENCE
*Articles in Advance*, pp. 1–20
ISSN 0025-1909 (print), ISSN 1526-5501 (online)

# Robust Optimization with Decision-Dependent Information Discovery

**Phebe Vayanos,[a,b,c,*] Angelos Georghiou,[d] Han Yu[b]**

[a] Center for Artificial Intelligence in Society, University of Southern California, Los Angeles, California 90089; [b] Daniel J. Epstein Department of Industrial & Systems Engineering, Viterbi School of Engineering, University of Southern California, Los Angeles, California 90089; [c] Thomas Lord Department of Computer Science, Viterbi School of Engineering, University of Southern California, Los Angeles, California 90089; [d] Department of Business and Public Administration, University of Cyprus, 1678 Nicosia, Cyprus

*Corresponding author

**Contact:** phebe.vayanos@usc.edu, https://orcid.org/0000-0001-7800-7235 (PV); georghiou.angelos@ucy.ac.cy, https://orcid.org/0000-0003-4490-4020 (AG); hyu376@usc.edu (HY)

**Received:** January 28, 2021
**Revised:** August 22, 2022; July 9, 2024
**Accepted:** September 10, 2024
**Published Online in Articles in Advance:** June 20, 2025

https://doi.org/10.1287/mnsc.2021.00160



**Abstract.** Robust optimization (RO) is a popular paradigm for modeling and solving two- and multistage decision-making problems affected by uncertainty. In many real-world applications, such as R&D project selection, production planning, or preference elicitation for product or policy recommendations, the *time of information discovery* is decision-dependent and the uncertain parameters only become observable after an often costly investment. Yet, most of the literature on robust optimization assumes that the uncertain parameters can be observed *for free* and that the sequence in which they are revealed is *independent* of the decision-maker's actions. To fill this gap in the practicability of RO, we consider two- and multistage robust optimization problems in which part of the decision variables control the time of information discovery. Thus, *information* available at any given time is *decision-dependent* and can be discovered (at least in part) by making strategic exploratory investments in previous stages. We propose a novel *dynamic* formulation of the problem and prove its correctness. We leverage our model to provide a solution method inspired from the $K$-adaptability approximation, whereby $K$ candidate strategies for each decision stage are chosen *here-and-now* and, at the beginning of each period, the best of these strategies is selected *after* the uncertain parameters that were chosen to be observed are revealed. We reformulate the problem as a finite mixed-integer (resp. bilinear) program if none (resp. some) of the decision variables are real-valued. This finite program is solvable with off-the-shelf solvers. We generalize our approach to the minimization of piecewise linear convex functions. We demonstrate the effectiveness of our method in terms of usability, optimality, and speed on synthetic instances of the Pandora box problem, the preference elicitation problem with real-valued recommendations, the best box problem, and the R&D project portfolio optimization problem. Finally, we evaluate it on an instance of the active preference elicitation problem used to recommend kidney allocation policies to policy-makers at the United Network for Organ Sharing based on real data from the U.S. Kidney Allocation System.

**History:** Accepted by Chung Piaw Teo, optimization.
**Funding:** This work was supported primarily by the Operations Engineering Program of the National Science Foundation under NSF Award No. 1763108. The authors are grateful for this support.
**Supplemental Material:** The online appendix and data files are available at https://doi.org/10.1287/mnsc.2021.00160.

**Keywords:** robust optimization • endogenous uncertainty • decision-dependent information discovery • Pandora box problem • R&D project portfolio selection • preference elicitation • kidney allocation

## 1. Introduction

### 1.1. Background & Motivation

Over the last two decades, robust optimization has emerged as a popular approach for decision-making under uncertainty in both *single-* and *multistage* settings, see for example, Ben-Tal et al. (2009), Ben-Tal and Nemirovski (2000), Ben-Tal and Nemirovski (1999), Ben-Tal and Nemirovski (1998), Bertsimas et al. (2004), Bertsimas and Sim (2004), Ben-Tal et al. (2004), Bertsimas et al. (2011), Zhen et al. (2018), Vayanos et al. (2012), Bertsimas and Goyal (2012), and Xu and Burer (2018). In multistage models, the uncertain parameters are revealed sequentially as time progresses and the decisions are allowed to depend on all the information made available in the past. Mathematically, decisions are modeled as functions of the history of observations, thus capturing the *adaptive* and *nonanticipative* nature of the decision process.

Most models and solution approaches in multistage robust optimization are tailored to problems where the uncertain parameters are *exogenous*, being independent of the decision-maker's actions. In particular, they assume that uncertainties can be observed *for free* and that the *sequence* in which they are revealed *cannot be influenced* by the decision-maker. Yet, these assumptions fail to hold in many real-world applications where the *time of information discovery* is decision-dependent and the uncertain parameters only become observable after an often costly investment. Mathematically, some binary *measurement* (or *observation*) decisions control the time of information discovery and the nonanticipativity requirements depend upon these decisions, severely complicating solution.

We now detail several applications areas where the time of revelation of the uncertain parameters is decision-dependent.

**R&D Project Portfolio Optimization.** Research and development firms typically maintain long pipelines of candidate projects whose returns are uncertain, see Solak et al. (2010). For each project, the firm can decide whether and when to start it and the amount of resources to be allocated to it. The return of each project will only be revealed once the project is completed. Thus, project start times and resource allocation decisions impact the time of information discovery in this problem.

**Clinical Trial Planning.** Pharmaceutical companies typically maintain long R&D pipelines of candidate drugs, see for example, Colvin and Maravelias (2008). Before any drug can reach the marketplace it needs to pass a number of costly clinical trials whose outcome (success/failure) is uncertain and will only be revealed after the trial is completed. Thus, the decisions to proceed with a trial control the time of information discovery in this problem.

**Offshore Oilfield Exploitation.** Offshore oilfields consist of several reservoirs of oil whose volume and initial deliverability (maximum initial extraction rate) are uncertain, see for example, Jonsbråten (1998), Goel and Grossman (2004), and Vayanos et al. (2011). While seismic surveys can help estimate these parameters, current technology is not sufficiently advanced to obtain accurate estimates. In fact, the volume and deliverability of each reservoir only become precisely known if a very expensive oil platform is built at the site and the drilling process is initiated. Thus, the decisions to build a platform and drill into a reservoir control the time of information discovery in this problem.

**Production Planning.** Manufacturing companies can typically produce a large number of different items. For each type of item, they can decide whether and how much to produce to satisfy their demand given that certain items are substitutable, see for example, Jonsbråten et al. (1998). The production cost of each item type is unknown and will only be revealed if the company chooses to produce the item. Thus, the decisions to produce a particular type of item control the time of information discovery in this problem.

**Active Preference Elicitation.** Preference elicitation refers to the problem of developing a decision support system capable of generating recommendations to a user, thus assisting in decision making. In active preference elicitation, one can ask users a (typically limited) number of questions from a potentially large set before making a recommendation, see for example, Vayanos et al. (2022). The answers to the questions are initially unknown and will only be revealed if the particular question is asked. Thus, the choices of questions to ask control the time of information discovery in this problem.

### 1.2. Literature Review

**Decision-Dependent Information Discovery.** Our paper relates to research on optimization problems affected by uncertain parameters whose time of revelation is decision-dependent and which originates in the literature on stochastic programming. The vast majority of these works assumes that the uncertain parameters are discretely distributed. In such cases, the decision process can be modeled by means of a finite scenario tree whose branching structure depends on the binary measurement decisions that determine the time of information discovery. This research began with the works of Jonsbråten et al. (1998) and Jonsbråten (1998). Jonsbråten et al. (1998) consider the case where all measurement decisions are made in the first stage and propose a solution approach based on an implicit enumeration algorithm. Jonsbråten (1998) generalizes this enumeration-based framework to the case where measurement decisions are made over time. More recently, Goel and Grossman (2004) showed that stochastic programs with discretely distributed uncertain parameters whose time of revelation is decision-dependent can be formulated as deterministic mixed-binary programs whose size is exponential in the number of endogenous uncertain parameters. To help deal with the "curse of dimensionality," they propose to precommit all measurement decisions, for example, to approximate them by here-and-now decisions, and to solve the multistage problem using either a decomposition technique or a folding horizon approach. Later, Goel and Grossman (2006), Goel et al. (2006), and Colvin and Maravelias (2010) propose optimization-based solution techniques that truly account for the adaptive nature of the measurement decisions and that rely on branch-and-bound and branch-and-cut approaches, respectively. Accordingly, Colvin and Maravelias (2010) and Gupta and Grossmann (2011) have proposed iterative solution schemes based on relaxations of the nonanticipativity constraints for the measurement variables. Our paper most closely relates to the work of Vayanos et al. (2011), wherein the authors investigate two- and multistage stochastic and robust programs with decision-dependent information discovery that involve continuously distributed uncertain parameters. They propose a decision-rule based approximation approach that relies on a prepartitioning of the support of the uncertain parameters. Since this approach applies in our context, we will benchmark against it in our experiments.

**Robust Optimization with Decision-Dependent Uncertainty Sets.** Our work also relates to the literature on robust optimization with uncertainty sets parameterized by the decisions. Such problems capture the ability of the decision-maker to influence the set of possible realizations of the uncertain parameters and have been investigated by Spacey et al. (2012), Nohadani and Sharma (2018), Nohadani and Roy (2017), Zhang et al. (2017), and Bertsimas and Vayanos (2017). These models do not apply in our context since they do not capture the ability of the decision-maker to influence the *information* available. In particular, the problems investigated by Spacey et al. (2012), Nohadani and Sharma (2018), and Nohadani and Roy (2017) are all single-stage, while *problems with decision-dependent information discovery are inherently sequential in nature*.

**Distributionally Robust Optimization with Decision-Dependent Ambiguity Sets.** Similarly, our paper is related to research on distributionally robust optimization with decision-dependent ambiguity sets, see for example, Luo and Mehrotra (2020), Basciftci et al. (2021), Noyan et al. (2022), and Yu and Shen (2022). To the best of our knowledge none of these works consider the case of decision-dependent information discovery which is the focus of our work.

**Robust Optimization with Binary Adaptive Variables.** Two-stage, and to a lesser extent also multistage, robust binary optimization problems have received considerable attention in the recent years. One stream of works proposes to restrict the functional form of the recourse decisions to functions of benign complexity, see Bertsimas and Dunn (2017) and Bertsimas and Georghiou (2015, 2018). A second stream of work relies on partitioning the uncertainty set into finite sets and applying constant decision rules on each partition, see Vayanos et al. (2011), Bertsimas and Dunning (2016), Postek and Den Hertog (2016), Bertsimas and Vayanos (2017). The last stream of works investigates the so-called $K$-adaptability counterpart of two-stage problems, see Bertsimas and Caramanis (2010), Hanasusanto et al. (2015), Subramanyam et al. (2020), Chassein et al. (2019), and Rahmattalabi et al. (2019). In this approach, $K$ candidate policies are chosen here-and-now and the best of these policies is selected after the uncertain parameters are revealed. Most of these papers assume that the uncertain parameters are *exogenous* in the sense that they are *in*dependent of the decision-maker's actions. Our paper most closely relates to the works of Bertsimas and Caramanis (2010) and Hanasusanto et al. (2015). It generalizes and subsumes the approach from Hanasusanto et al. (2015) to problems with *decision-dependent information discovery*, to *multistage* problems, and to problems with *piecewise linear convex objective*.

**Stochastic Probing.** Our paper also fits in a line of work on stochastic probing in the computer science literature, see Gupta et al. (2016, 2017) and Singla (2018). Here, the problem consists of a set of elements with uncertain value whose distribution is known but whose realization becomes observable only after the element is probed. However, probing is costly (incurs a cost or consumes budget) and irrevocable and the goal is to choose the set of elements to probe and the order in which to probe them to maximize profit (e.g., the value of the item with the highest value that has been probed). Concrete examples include the best box problem and the Pandora box problem, see for example, Singla (2018). The techniques presented in this stream of work do not apply to the case where the distributions are unknown, to general optimization problems with decision-dependent information discovery, nor to problems with general, potentially uncertain, constraints.

**Worst-Case Regret Optimization.** Finally, our work relates to two-stage worst-case absolute regret minimization problems, see for example, Assavapokee et al. (2008a, b), Zhang (2011), Jiang et al. (2013), Ng (2013), Chen et al. (2014), Ning and You (2018), and Poursoltani and Delage (2019). To the best of our knowledge, our paper is the first to investigate worst-case regret minimization problems in the presence of uncertain parameters whose time of revelation is decision-dependent.

### 1.3. Proposed Approach and Contributions

We now summarize our approach and main contributions in this paper:

a. We consider general two- and multistage robust optimization problems with decision-dependent information discovery. These encompass as special cases the R&D project portfolio optimization problem, the Pandora box problem (which can be used to model job candidate selection and house hunting, among others), the active preference elicitation problem, and many more. To the best of our knowledge, only one other paper in the literature studies such problems in the robust optimization setting. We propose novel "min-max-min-max-⋯-min-max" reformulations of these problems and prove correctness of our formulations. These reformulations unlock new approximate (and potentially also exact) solution approaches for addressing problems with decision-dependent information discovery.

b. We leverage our new reformulations to propose a solution approach based on the $K$-adaptability approximation, wherein $K$ candidate strategies are chosen here-and-now and the best of these strategies is selected after the uncertain parameters that were chosen to be observed are revealed. This approximation allows us to control the trade-off between complexity and solution quality by tuning a single design parameter, $K$. We propose practicable reformulations of the $K$-adaptability counterpart of problems with decision-dependent information discovery in the form of moderately sized finite programs solvable with off-the shelf solvers. These programs can be written equivalently as mixed-binary linear programs if all decision-variables are binary. Our reformulations subsume those from the literature that apply only to two-stage problems with exogenous uncertain parameters.

c. We generalize the $K$-adaptability approximation scheme to multistage problems and to problems with piecewise linear convex objective function. The piecewise linear convex objective enables us, among others, to address worst-case absolute regret minimization problems. These generalizations and associated algorithm that we provide apply also to problems with exogenous uncertain parameters.

d. We perform a wide array of experiments on the R&D project portfolio selection problem, the preference elicitation problem with real-valued recommendations, the best box selection problem, Pandora's box problem, and the preference elicitation problem. We show that our proposed approach outperforms the state-of-the-art in the literature in terms of usability, optimality, and speed. Indeed, our approach reduces the number of subsets in the recourse strategy by a factor of 3, improves the quality of the returned solution by a factor of 1.9, and results in an 8.5× speed-up. We perform a case study showcasing the benefits of our approach on real data from the U.S. Kidney Allocation System (KAS) to recommend policies that meet the needs of policy-makers at the Organ Procurement and Transplantation Network (OPTN) and the United Network for Organ Sharing (UNOS), the lead agency in charge of allocating organs for transplantation in the United States.[1]

### 1.4. Organization of the Paper and Notation

The paper is organized as follows. Sections 2 and 3 introduce two-stage robust optimization problems with exogenous uncertainty and with decision-dependent information discovery (DDID), respectively. In particular, Section 3 introduces our novel formulation. Section 4 proposes reformulations of the $K$-adaptability counterparts of problems with DDID as finite programs solvable with off-the-shelf solvers. Section 5 generalizes the $K$-adaptability approximation to problems with piecewise linear convex objective and proposes an efficient solution procedure. Section 6 presents computational results on synthetic instances of the two-stage best box selection problem, the two-stage preference elicitation problem with real-valued recommendations, and the two-stage R&D project portfolio optimization problem. Finally, Section 7 formulates the preference elicitation problem for learning the preferences of policy-makers at the OPTN/UNOS as a two-stage robust problem with decision-dependent information discovery, and presents numerical results on real data from the U.S. Kidney Allocation System. The proofs of all statements can be found in the Electronic Companion to the paper. Proposed extensions to our methods, algorithms, and speed-up strategies are also deferred to the Electronic Companion. In particular, Sections EC.1 and EC.2 generalize the $K$-adaptability approximation to multistage problems and apply it to the multistage Pandora's box problem, respectively.

**Notation.** Throughout this paper, vectors (matrices) are denoted by boldface lowercase (uppercase) letters. The $k$th element of a vector $\boldsymbol{x} \in \mathbb{R}^n$ ($k \le n$) is denoted by $\boldsymbol{x}_k$. Scalars are denoted by lowercase letters, for example, $\alpha$ or $u$. For a matrix $\boldsymbol{H} \in \mathbb{R}^{n \times m}$, we let $[\boldsymbol{H}]_k \in \mathbb{R}^m$ denote the $k$th row of $\boldsymbol{H}$, written as a column vector. We let $\mathcal{L}_n^k$ denote the space of all functions from $\mathbb{R}^n$ to $\mathbb{R}^k$. Accordingly, we denote by $\mathcal{B}_n^k$ the spaces of all functions from $\mathbb{R}^n$ to $\{0, 1\}^k$. Given two vectors of equal length, $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^n$, we let $\boldsymbol{x} \circ \boldsymbol{y}$ denote the Hadamard product of the vectors, for example, their elementwise product. Given a set $\mathcal{A}$ and a positive integer $n$, we let $\mathcal{A}^n := \mathcal{A} \times \mathcal{A} \times \cdots \times \mathcal{A}$ ($n$ times). With a slight abuse of notation, we may use the maximum and minimum operators even when the optimum may not be attained; in such cases, the operators should be understood as suprema and infima, respectively. We use the convention that a decision is feasible for a minimization problem if and only if it attains an objective that is $< +\infty$. Finally, for a logical expression $E$, we define the indicator function $\mathbb{I}(E)$ as $\mathbb{I}(E) := 1$ if $E$ is true and 0 otherwise.

## 2. Two-Stage RO with Exogenous Uncertainty

To motivate our formulation from Section 3, we introduce two equivalent models of two-stage robust optimization with *exogenous uncertainty* from the literature and discuss their relative merits.

In two-stage robust optimization with *exogenous* uncertainty, first-stage (or here-and-now) decisions $\boldsymbol{x} \in \mathcal{X} \subseteq \mathbb{R}^{N_x}$ are made today, *before* any of the uncertain parameters are observed. Subsequently, all of the uncertain parameters $\boldsymbol{\xi} \in \Xi \subseteq \mathbb{R}^{N_\xi}$ are revealed. Finally, once the realization of $\boldsymbol{\xi}$ has become available, second-stage (or wait-and-see) decisions $\boldsymbol{y} \in \mathcal{Y} \subseteq \mathbb{R}^{N_y}$ are selected. We assume that the uncertainty set $\Xi$ is a nonempty bounded polyhedron expressible as $\Xi := \{\boldsymbol{\xi} \in \mathbb{R}^{N_\xi} : \boldsymbol{A\xi} \le \boldsymbol{b}\}$ for some matrix $\boldsymbol{A} \in \mathbb{R}^{R \times N_\xi}$ and vector $\boldsymbol{b} \in \mathbb{R}^R$. As the decisions $\boldsymbol{y}$ are selected after the uncertain parameters are revealed, they are allowed to *adapt* or *adjust* to the realization of $\boldsymbol{\xi}$. In the literature, there are two formulations of generic two-stage robust problem with exogenous uncertainty: they differ in the way in which the ability of $\boldsymbol{y}$ to adapt to $\boldsymbol{\xi}$ is modeled.

### 2.1. Decision Rule Formulation

In the first model, one optimizes today over both the here-and-now decisions $\boldsymbol{x}$ and over recourse actions $\boldsymbol{y}$ to be taken in each realization of $\boldsymbol{\xi}$. The decision $\boldsymbol{y}$ is modeled as a function (or *decision rule*) of $\boldsymbol{\xi}$ that is selected today, along with $\boldsymbol{x}$. Under this paradigm, a two-stage linear robust problem with exogenous uncertainty is expressible as:

$$\begin{array}{ll} \text{minimize} & \max\limits_{\boldsymbol{\xi} \in \Xi} \; \boldsymbol{\xi}^\top \boldsymbol{C}\boldsymbol{x} + \boldsymbol{\xi}^\top \boldsymbol{Q}\boldsymbol{y}(\boldsymbol{\xi}) \\ \text{subject to} & \boldsymbol{x} \in \mathcal{X},\, \boldsymbol{y} \in \mathcal{L}_{N_\xi}^{N_y} \\ & \left.\begin{array}{l} \boldsymbol{y}(\boldsymbol{\xi}) \in \mathcal{Y} \\ \boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}(\boldsymbol{\xi}) \le \boldsymbol{H}\boldsymbol{\xi} \end{array}\right\} \quad \forall \boldsymbol{\xi} \in \Xi, \end{array} \tag{1}$$

where $\boldsymbol{C} \in \mathbb{R}^{N_\xi \times N_x}$, $\boldsymbol{Q} \in \mathbb{R}^{N_\xi \times N_y}$, $\boldsymbol{T} \in \mathbb{R}^{L \times N_x}$, $\boldsymbol{W} \in \mathbb{R}^{L \times N_y}$, and $\boldsymbol{H} \in \mathbb{R}^{L \times N_\xi}$. We assume that the objective function and right hand-sides are *linear* in $\boldsymbol{\xi}$. We can account for *affine* dependencies on $\boldsymbol{\xi}$ by introducing an auxiliary uncertain parameter $\boldsymbol{\xi}_{N_\xi + 1}$ restricted to equal unity.

### 2.2. Min-Max-Min Formulation

In the second model, only $\boldsymbol{x}$ is selected today and the recourse decisions $\boldsymbol{y}$ are optimized explicitly, in a *dynamic* fashion, *after* nature is done making a decision. Under this model, a two-stage robust problem with exogenous uncertainty is expressible as:

$$\begin{array}{ll} \text{minimize} & \max\limits_{\boldsymbol{\xi} \in \Xi} \left[ \boldsymbol{\xi}^\top \boldsymbol{C}\boldsymbol{x} + \min\limits_{\boldsymbol{y} \in \mathcal{Y}} \{ \boldsymbol{\xi}^\top \boldsymbol{Q}\boldsymbol{y} : \boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y} \le \boldsymbol{H}\boldsymbol{\xi} \} \right] \\ \text{subject to} & \boldsymbol{x} \in \mathcal{X}. \end{array} \tag{2}$$

Problems (1) and (2) are equivalent, see for example, Shapiro (2017). However, each of them has proved successful in different contexts. Problem (1) has been the building block of most of the literature on the decision rule approximation, see Section 1. Problem (2) has enabled the advent and tremendous success of the $K$-adaptability approximation approach to two-stage robust problems with binary recourse, see Bertsimas and Caramanis (2010), Hanasusanto et al. (2015). It has also facilitated the development of algorithms and efficient solution schemes, see for example, Zeng and Zhao (2013), Ayoub and Poss (2016), and Bertsimas and Shtern (2018).

## 3. Two-Stage RO with Decision-Dependent Information Discovery

In this section, we describe two-stage robust optimization problems with decision-dependent information discovery (DDID) and propose an entirely new modeling framework for studying such problems. This framework underpins our ability to generalize the popular $K$-adaptability approximation approach from the literature to problems affected by uncertain parameters whose time of revelation is decision-dependent, see Sections 4.2 and 4.3.

### 3.1. Problem Description

In two-stage robust optimization with DDID, the uncertain parameters $\boldsymbol{\xi}$ do not necessarily become observed (for free) between the first and second decision-stages. Instead, some (typically costly) first stage decisions control the *time of information discovery* in the problem: they decide whether (and which of) the uncertain parameters will be revealed *before* the wait-and-see decisions $\boldsymbol{y}$ are selected. If the decision-maker chooses to not observe some of the uncertain parameters, then those parameters will still be uncertain at the time when the decision $\boldsymbol{y}$ is selected, and $\boldsymbol{y}$ will only be allowed to depend on the portion of the uncertain parameters that have been revealed. On the other hand, if the decision-maker chooses to observe all of the uncertain parameters, then there will be no uncertainty in the problem at the time when $\boldsymbol{y}$ is selected, and $\boldsymbol{y}$ will be allowed to depend on all uncertain parameters.

In order to allow for endogenous uncertainty, we introduce a here-and-now binary measurement (or observation) decision vector $\boldsymbol{w} \in \{0,1\}^{N_\xi}$ of the same dimension as $\boldsymbol{\xi}$ whose $i$th element $\boldsymbol{w}_i$ is 1 if and only if we choose to observe $\boldsymbol{\xi}_i$ between the first and second decision stages. In the presence of such endogenous uncertain parameters, the recourse decisions $\boldsymbol{y}$ are selected after the *portion* of uncertain parameters that was *chosen* to be observed is revealed. In particular, $\boldsymbol{y}$ must be constant in (i.e., robust to) those uncertain parameters that remain unobserved at the second decision-stage. The requirement that $\boldsymbol{y}$ only depend on the uncertain parameters that have been revealed at the time it is chosen is termed *nonanticipativity*. In the presence of uncertain parameters whose time of revelation is decision-dependent, this requirement translates to *decision-dependent nonanticipativity constraints*.

### 3.2. Decision Rule Formulation

In the literature and to the best of our knowledge, two-stage robust optimization problems with DDID have been formulated (in a manner paralleling Problem (1)) by letting the recourse decisions $\boldsymbol{y}$ be functions of $\boldsymbol{\xi}$ and requiring that those functions be constant in $\boldsymbol{\xi}_i$ if $w_i = 0$, see Vayanos et al. (2011). Under this (decision rule based) modeling paradigm, generic two-stage robust optimization problems with decision-dependent information discovery take the form

$$\begin{array}{ll} \text{minimize} & \max_{\boldsymbol{\xi}\in\Xi} \; \boldsymbol{\xi}^\top \boldsymbol{C}\boldsymbol{x} + \boldsymbol{\xi}^\top \boldsymbol{D}\boldsymbol{w} + \boldsymbol{\xi}^\top \boldsymbol{Q}\boldsymbol{y}(\boldsymbol{\xi}) \\ \text{subject to} & \boldsymbol{x}\in\mathcal{X},\, \boldsymbol{w}\in\mathcal{W},\, \mathbf{y}\in\mathcal{L}_{N_\xi}^{N_y} \\ & \left.\begin{array}{l} \boldsymbol{y}(\boldsymbol{\xi})\in\mathcal{Y} \\ \boldsymbol{T}\boldsymbol{x}+\boldsymbol{V}\boldsymbol{w}+\boldsymbol{W}\boldsymbol{y}(\boldsymbol{\xi})\le\boldsymbol{H}\boldsymbol{\xi} \end{array}\right\} \quad \forall \boldsymbol{\xi}\in\Xi \\ & \boldsymbol{y}(\boldsymbol{\xi}) = \boldsymbol{y}(\boldsymbol{\xi}') \quad \forall \boldsymbol{\xi},\boldsymbol{\xi}'\in\Xi : \boldsymbol{w}\circ\boldsymbol{\xi} = \boldsymbol{w}\circ\boldsymbol{\xi}', \end{array} \tag{3}$$

where $\mathcal{W}\subseteq\{0,1\}^{N_\xi}$, $\boldsymbol{D}\in\mathbb{R}^{N_\xi\times N_\xi}$, $\boldsymbol{V}\in\mathbb{R}^{L\times N_\xi}$, and the remaining data elements are as in Problem (1). The set $\mathcal{W}$ can encode requirements on the measurement decisions. For example, it can enforce that a given uncertain parameter $\boldsymbol{\xi}_i$ may only be observed if another uncertain parameter $\boldsymbol{\xi}_{i'}$ has been observed using $w_i \le w_{i'}$. Accordingly, it can postulate that the total number of uncertain parameters that are observed does not exceed a certain budget $Q$ using $\sum_{i=1}^{N_\xi} w_i \le Q$. If only some (or all) of the uncertain parameters have a time of information discovery that is exogenous, our models and solution approaches can be used by restricting the observation decisions $w_i$ to equal 1 (resp. 0) for each *exogenous* uncertain parameter $i$ that is (resp. is not) observed between the first and second decision stages. These restrictions can be conveniently added as constraints to the set $\mathcal{W}$. The last constraint in the problem is a decision-dependent nonanticipativity constraint: it ensures that the function $\boldsymbol{y}$ is constant in the uncertain parameters that remain unobserved at the second stage. Indeed, the identity $\boldsymbol{w}\circ\boldsymbol{\xi} = \boldsymbol{w}\circ\boldsymbol{\xi}'$ evaluates to true only if the elements of $\boldsymbol{\xi}$ and $\boldsymbol{\xi}'$ that were observed are indistinguishable, in which case the decisions taken in scenarios $\boldsymbol{\xi}$ and $\boldsymbol{\xi}'$ must be equal. We omit joint (first stage) constraints on $\boldsymbol{x}$ and $\boldsymbol{w}$ to minimize notational overhead but emphasize that our approach remains applicable in their presence.

Note that Problem (3) generalizes Problem (1). Indeed, if we set $\boldsymbol{w} = \mathbf{e}$, $\boldsymbol{D} = \mathbf{0}$, and $\boldsymbol{V} = \mathbf{0}$ in Problem (3), we recover Problem (1). In addition, it generalizes the single-stage robust problem: if we set $\boldsymbol{w} = \mathbf{0}$ in Problem (3), all uncertain parameters are revealed *after* the second stage so that the second stage decisions are forced to be static (i.e., constant in $\boldsymbol{\xi}$).

To the best of our knowledge, the only approach in the literature for (approximately) solving problems of type (3) is presented in Vayanos et al. (2011) and relies on a decision rule approximation. The authors propose to approximate the binary (resp. continuous) wait-and-see decisions by functions that are piecewise constant (resp. piecewise linear) on a preselected partition of the uncertainty set of the form $\Xi_{\boldsymbol{s}} := \{\boldsymbol{\xi}\in\Xi : \boldsymbol{c}^i_{s_i-1} \le \boldsymbol{\xi}_i < \boldsymbol{c}^i_{s_i},\, i=1,\ldots,k\}$, where $\boldsymbol{s}\in\mathcal{S} := \times_{i=1}^{N_\xi}\{1,\ldots,\boldsymbol{r}_i\}\subseteq\mathbb{Z}^{N_\xi}$ and $\boldsymbol{c}^i_1 < \boldsymbol{c}^i_2 < \cdots < \boldsymbol{c}^i_{r_i-1}$ for $i=1,\ldots,N_\xi$ represent $\boldsymbol{r}_i - 1$ breakpoints along the $\boldsymbol{\xi}_i$ axis. Unfortunately, as the following example illustrates, this approach is highly sensitive to the choice of breakpoint configuration.

**Example 1.** Consider the following instance of Problem (3)

$$\begin{array}{ll} \text{minimize} & 0 \\ \text{subject to} & \boldsymbol{w}\in\{0,1\}^2,\, \boldsymbol{y}\in\mathcal{B}_2^2 \\ & \left.\boldsymbol{\xi}-\boldsymbol{\epsilon} \le \boldsymbol{y}(\boldsymbol{\xi}) \le \mathbf{e}+\boldsymbol{\xi}-\boldsymbol{\epsilon}\right\} \quad \forall\boldsymbol{\xi}\in\Xi \\ & \boldsymbol{y}(\boldsymbol{\xi}) = \boldsymbol{y}(\boldsymbol{\xi}') \quad \forall\boldsymbol{\xi},\boldsymbol{\xi}'\in\Xi : \boldsymbol{w}\circ\boldsymbol{\xi} = \boldsymbol{w}\circ\boldsymbol{\xi}', \end{array} \tag{4}$$

where $\Xi := [-1,1]^2$. The inequality constraints in the problem combined with the requirement that $\boldsymbol{y}(\boldsymbol{\xi})$ be binary imply that we must have $\boldsymbol{y}_i(\boldsymbol{\xi}) = 1$ (resp. 0) whenever $\boldsymbol{\xi}_i > \epsilon_i$ (resp. $\boldsymbol{\xi}_i < \epsilon_i$). Thus, from the decision-dependent nonanticipativity constraints, the only feasible choice for $\boldsymbol{w}$ is $\mathbf{e}$. It is easy to show that if $\boldsymbol{\epsilon} = 1\text{e}-3\mathbf{e}$ and if we uniformly partition each axis iteratively in 2, 3, 4, etc. subsets, then 1999 breakpoints along each direction will need to be introduced before reaching a feasible (and thus optimal) solution. The associated problem will involve over 8e7 binary decision variables and 16e7 constraints. In contrast, as will become clear later on, our proposed solution approach with approximation parameter $K = 4$ will be optimal in this case.

Example 1 is not surprising: the approach from Vayanos et al. (2011) was motivated by stochastic programs which are less sensitive to the breakpoint configuration than robust problems. Thus, a more flexible approach is needed to address two-stage and multistage robust problems with DDID.

### 3.3. Proposed Min-Max-Min-Max Formulation

Motivated by the success of formulation (2) as the starting point to solve two-stage robust optimization problems with exogenous uncertainty, we derive an analogous *dynamic* formulation for the case of endogenous uncertainties. In particular, we build a robust optimization problem in which the sequence of problems solved by each of the decision-maker and nature in turn is captured explicitly. The idea is as follows. Initially, the decision-maker selects $\boldsymbol{x}\in\mathcal{X}$ and $\boldsymbol{w}\in\mathcal{W}$. Subsequently, nature commits to a realization $\overline{\boldsymbol{\xi}}$ of the uncertain parameters from the set $\Xi$. Then, the decision-maker selects a recourse action $\boldsymbol{y}$ that needs to be robust to those elements $\overline{\boldsymbol{\xi}}_i$ of the uncertain vector $\overline{\boldsymbol{\xi}}$ that they have not observed, that is, for which $w_i = 0$. Indeed, the decision $y$ may have to be taken under uncertainty if there is some $i$ such that $w_i = 0$, in which case not all of the uncertain parameters have been revealed when $y$ is selected. Indeed, after $y$ is selected, nature is free to choose any realization of $\xi \in \Xi$ that is compatible with the original choice $\overline{\xi}$ in the sense that $\xi_i = \overline{\xi}_i$ for all $i$ such that $w_i = 1$. This model captures the notion that, after $y$ has been selected, nature is still free to choose the elements $\xi_i$ that have not been observed, provided it does so in a way that is consistent with those parameters that *have* been observed. Mathematically, given the measurement decisions $w$ and the observation $\overline{\xi}$, nature can select any element $\xi$ from the set

$$\Xi(w, \overline{\xi}) := \{\xi \in \Xi : w \circ \xi = w \circ \overline{\xi}\}.$$

Note in particular that if $w = \mathbf{e}$, then $\Xi(w, \overline{\xi}) = \{\overline{\xi}\}$ and there is no uncertainty when $y$ is chosen. Accordingly, if $w = \mathbf{0}$, then $\Xi(w, \overline{\xi}) = \Xi$ and $y$ has no knowledge of any of the elements of $\xi$. The realizations $\overline{\xi}$, $\xi$, and the sets $\Xi$ and $\Xi(w, \overline{\xi})$ are all illustrated on Figure 1.

Based on the above notation, we propose the following generic formulation of a two-stage robust optimization problem with decision-dependent information discovery:

$$\begin{aligned} \min\ \max_{\overline{\xi} \in \Xi}\ \min_{y \in \mathcal{Y}} & \left\{ \max_{\xi \in \Xi(w, \overline{\xi})} \xi^\top C x + \xi^\top D w + \xi^\top Q y \,:\, Tx \right. \\ & \left. + Vw + Wy \le H\xi \quad \forall \xi \in \Xi(w, \overline{\xi}) \right\} \qquad (\mathcal{P}) \\ \text{s.t.}\ & x \in \mathcal{X},\ w \in \mathcal{W}. \end{aligned}$$

Note that, at the time when $y$ is selected, some elements of $\xi$ are still uncertain. The choice of $y$ thus needs to be robust to the choice of those uncertain parameters that remain to be revealed. In particular, the constraints need to be satisfied for all choices of $\xi \in \Xi(w, \overline{\xi})$. Accordingly, $y$ is chosen so as to minimize the worst-case possible cost when $\xi$ is valued in the set $\xi \in \Xi(w, \overline{\xi})$.

Problems (3) and ($\mathcal{P}$) are equivalent in a sense made precise in the following theorem.

**Theorem 1.** *The optimal objective values of Problems (3) and ($\mathcal{P}$) are equal. Moreover, the following statements hold true:*

i. *Let $(x, w)$ be optimal in ($\mathcal{P}$) and, for each $\delta$ such that $\delta = w \circ \overline{\xi}$ for some $\overline{\xi} \in \Xi$, define*

$$\begin{aligned} y'(\delta) \in \arg\min_{y \in \mathcal{Y}} & \left\{ \max_{\xi \in \Xi(w, \delta)} \xi^\top C x + \xi^\top D w + \xi^\top Q y \,:\, Tx \right. \\ & \left. + Vw + Wy \le H\xi \quad \forall \xi \in \Xi(w, \delta) \right\}. \end{aligned}$$

*Also, for each $\xi \in \Xi$, define $y(\xi) := y'(w \circ \xi)$. Then, $(x, w, y(\cdot))$ is optimal in Problem (3).*

ii. *Let $(x, w, y(\cdot))$ be optimal in Problem (3). Then, $(x, w)$ is optimal in Problem ($\mathcal{P}$).*

The parameter $\delta$ in item (i) of the theorem above is introduced to ensure that the decision rule $y(\cdot)$ defined on $\Xi$ is nonanticipative. Indeed, if for any given $(x, w)$ and $\overline{\xi}$, there are many optimal solutions to problem

$$\begin{aligned} \min_{y \in \mathcal{Y}} & \left\{ \max_{\xi \in \Xi(w, \overline{\xi})} \xi^\top C x + \xi^\top D w + \xi^\top Q y \,:\, Tx + Vw \right. \\ & \left. + Wy \le H\xi \quad \forall \xi \in \Xi(w, \overline{\xi}) \right\}, \end{aligned}$$

the decision rule $\tilde{y}(\cdot)$ defined on $\Xi$ through

$$\begin{aligned} \tilde{y}(\overline{\xi}) \in \arg\min_{y \in \mathcal{Y}} & \left\{ \max_{\xi \in \Xi(w, \overline{\xi})} \xi^\top C x + \xi^\top D w + \xi^\top Q y \,:\, Tx \right. \\ & \left. + Vw + Wy \le H\xi \quad \forall \xi \in \Xi(w, \overline{\xi}) \right\}, \end{aligned}$$

may not be constant in those parameters that remain

**Figure 1.** (Color online) Companion Figure to Section 3.3

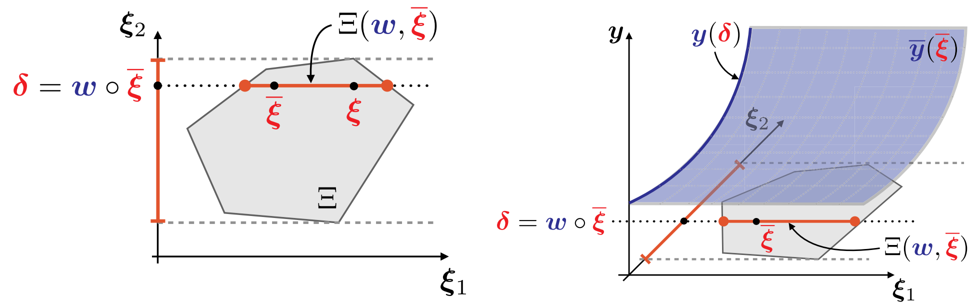


*Notes.* The figure on the left illustrates the role played by $\overline{\xi}$ in the new formulation ($\mathcal{P}$) and the definition of the uncertainty sets $\Xi$ and $\Xi(w, \overline{\xi})$. Consider a setting where $\Xi \subseteq \mathbb{R}^2$ (i.e., $N_\xi = 2$) and suppose that $w = (0, 1)$ so that the decision-maker has chosen to only observe $\xi_2$. In the figures, $\Xi$ is shown as the grey shaded area. Once $\overline{\xi}$ is chosen by nature, the decision-maker can only infer that $\xi$ will materialize in the set $\Xi(w, \overline{\xi})$ which collects all parameter realizations $\xi \in \Xi$ that satisfy $\xi_2 = \overline{\xi}_2$, being compatible with our partial observation. The figure on the right illustrates the construction of an optimal nonanticipative decision $\overline{y}$ from an optimal solution $y(\delta)$ to $\min_{y \in \mathcal{Y}} \{\max_{\xi \in \Xi(w, \delta)} \xi^\top C x + \xi^\top D w + \xi^\top Q y : Tx + Vw + Wy \le H\xi \quad \forall \xi \in \Xi(w, \delta)\}$, see Theorem 1. We note that the policy $\overline{y}$ constructed as in Theorem 1 is constant along the $\xi_1$ direction since here $w_1 = 0$.

unobserved. We note of course that other tie-breaking mechanisms could be used to build a nonanticipative solution. For example, we may select, among all optimal solutions, the one that is lexicographically first.

The theorem above is the main result that enables us to generalize the $K$-adaptability approximation scheme to two-stage robust problems with decision-dependent information discovery and binary recourse. In Electronic Companion EC.6, we show that for any given choice of here-and-now decisions, the set of parameters $\boldsymbol{\xi}$ for which a particular wait-and-see decision is optimal may be nonclosed and nonconvex and that the optimal value of the problem may not be attained. This result is expected from the analysis in Hanasusanto et al. (2015), since Problem ($\mathcal{P}$) generalizes Problem (2). Our example illustrates that this may be the case even if a portion of the uncertain parameters remain unobserved in the second stage.

Two-stage robust optimization problems with decision-dependent information discovery have a huge modeling power, see Sections 1, 6, and 7. Yet, as illustrated by the preceding discussion, they pose several theoretical and practical challenges. As we will see in the following sections, whether we are or not able to reformulate the $K$-adaptability counterpart of the problem exactly as a finite program solvable with off-the-shelf solvers depends on the absence or presence of uncertainty in the constraints. When in presence of constraint uncertainty, we can always compute an arbitrarily tight outer (lower bound) approximation, see Section 4.3.

## 4. *K*-Adaptability for Problems with DDID

Instead of solving Problem ($\mathcal{P}$) directly, we approximate it through its $K$-adaptability counterpart,

$$\begin{aligned} \min \ & \max_{\overline{\boldsymbol{\xi}} \in \Xi} \ \min_{k \in \mathcal{K}} \left\{ \max_{\boldsymbol{\xi} \in \Xi(\boldsymbol{w}, \overline{\boldsymbol{\xi}})} \ \boldsymbol{\xi}^\top \boldsymbol{C} \boldsymbol{x} + \boldsymbol{\xi}^\top \boldsymbol{D} \boldsymbol{w} + \boldsymbol{\xi}^\top \boldsymbol{Q} \boldsymbol{y}^k \ : \ \boldsymbol{T}\boldsymbol{x} \right. \\ & \qquad \left. + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{H}\boldsymbol{\xi} \quad \forall \boldsymbol{\xi} \in \Xi(\boldsymbol{w}, \overline{\boldsymbol{\xi}}) \right\} \qquad (\mathcal{P}_K) \\ \text{s.t.} \ & \boldsymbol{x} \in \mathcal{X}, \ \boldsymbol{w} \in \mathcal{W}, \ \boldsymbol{y}^k \in \mathcal{Y}, \ k \in \mathcal{K}, \end{aligned}$$

where $\mathcal{K} := \{1, \ldots, K\}$. In this problem, $K$ candidate policies $\boldsymbol{y}^1, \ldots, \boldsymbol{y}^K$ are chosen here-and-now, that is before $\boldsymbol{w} \circ \overline{\boldsymbol{\xi}}$ (the portion of uncertain parameters that we chose to observe) is revealed. Once $\boldsymbol{w} \circ \overline{\boldsymbol{\xi}}$ becomes known, the best of those policies among all those that are robustly feasible (in view of uncertainty in the uncertain parameters that are still unknown) is implemented. If all policies are infeasible for some $\overline{\boldsymbol{\xi}} \in \Xi$, then we interpret the maximum and minimum in ($\mathcal{P}_K$) as supremum and infimum, that is, the $K$-adaptability problem evaluates to $+\infty$. Problem ($\mathcal{P}_K$) is a conservative approximation to program ($\mathcal{P}$). Moreover, if $|\mathcal{Y}| < \infty$ and $K = |\mathcal{Y}|$, then the two problems are equivalent. In practice, we hope that a moderate number of candidate policies $K$ will be sufficient to obtain a (near) optimal solution to ($\mathcal{P}$).

### 4.1. The Price of Usability

Problem ($\mathcal{P}_K$) is interesting in its own right. Indeed, in problems where usability is important (e.g., if workers need to be trained to follow diverse contingency plans depending on the realization $\boldsymbol{w} \circ \overline{\boldsymbol{\xi}}$), Problem ($\mathcal{P}_K$) may be an attractive alternative to Problem ($\mathcal{P}$). In such settings, the loss in optimality incurred due to passing from Problem ($\mathcal{P}$) to Problem ($\mathcal{P}_K$) can be thought of as the *price of usability*. For example, consider an emergency response planning problem where, in the first stage, a small number of helicopters can be used to survey affected areas and, in the second stage, and in response to the observed state of the areas surveyed, deployment of emergency response teams is decided. In practice, to avoid having to train teams in a large number of plans (yielding significant operational challenges), only a moderate number of response plans may be allowed. The importance of interpretability/usability has been previously noted by for example, Koç and Morton (2015), McCarthy et al. (2018), Bertsimas et al. (2019), and Aghaei et al. (2019, 2024).

**Remark 1.** If $\mathcal{W} = \{0, 1\}^{N_\xi}$, $\boldsymbol{D} = \boldsymbol{0}$, and $\boldsymbol{V} = \boldsymbol{0}$, then $\boldsymbol{w} = \mathbf{e}$ is optimal in Problem ($\mathcal{P}$) and thus $\Xi(\boldsymbol{w}, \overline{\boldsymbol{\xi}}) = \{\overline{\boldsymbol{\xi}}\}$, implying that Problem ($\mathcal{P}$) reduces to Problem (2) and Problem ($\mathcal{P}_K$) reduces to the $K$-adaptability counterpart of Problem (2).

Relative to the problems studied by Bertsimas and Caramanis (2010) and Hanasusanto et al. (2015), Problem ($\mathcal{P}_K$) presents several challenges. First, the second stage problem in ($\mathcal{P}_K$) is a robust (as opposed to deterministic) optimization problem. Second, the uncertainty sets involved in the maximization tasks of this robust problem are decision-dependent. While Problem ($\mathcal{P}_K$) appears to be significantly more complicated than its exogenous counterpart, it can be converted to an equivalent min-max-min problem by *lifting* the space of the uncertainty set as show in the following lemma that is instrumental in our analysis.

**Lemma 1.** *The K-adaptability problem with decision-dependent information discovery, Problem ($\mathcal{P}_K$), is equivalent to*

$$\begin{aligned} \min \ & \max_{\{\boldsymbol{\xi}^k\}_{k \in \mathcal{K}} \in \Xi^K(\boldsymbol{w})} \ \min_{k \in \mathcal{K}} \{ (\boldsymbol{\xi}^k)^\top \boldsymbol{C} \boldsymbol{x} + (\boldsymbol{\xi}^k)^\top \boldsymbol{D} \boldsymbol{w} + (\boldsymbol{\xi}^k)^\top \boldsymbol{Q} \boldsymbol{y}^k \ : \\ & \qquad \boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{H}\boldsymbol{\xi}^k \} \qquad (5) \\ \text{s.t.} \ & \boldsymbol{x} \in \mathcal{X}, \ \boldsymbol{w} \in \mathcal{W}, \ \mathbf{y}^k \in \mathcal{Y}, \ k \in \mathcal{K}, \end{aligned}$$

*where*

$$\Xi^K(\boldsymbol{w}) := \{ \{\boldsymbol{\xi}^k\}_{k \in \mathcal{K}} \in \Xi^K \ : \ \exists \overline{\boldsymbol{\xi}} \in \Xi \ \text{such that} \ \boldsymbol{\xi}^k \in \Xi(\boldsymbol{w}, \overline{\boldsymbol{\xi}}) \ \text{for all} \ k \in \mathcal{K} \}. \qquad (6)$$

For any fixed $\boldsymbol{w} \in \mathcal{W}$, the subvector $\boldsymbol{\xi}^k$ in the definition of $\Xi^K(\boldsymbol{w})$ represents the uncertainty scenario that "nature" will choose if the decision-maker acts according to decisions $\boldsymbol{w}$ in the first stage and according to policy $k$ in the second stage. The set $\Xi^K(\boldsymbol{w})$ collects, for each $k \in \mathcal{K}$, all feasible choices that nature can take if the decision-maker acts according to $\boldsymbol{w}$ and then $\boldsymbol{y}^k$ in the first and second stages, respectively. Thus, in Problem (5), the decision-maker first selects $\boldsymbol{x}$, $\boldsymbol{w}$, and $\boldsymbol{y}^k$, $k \in \mathcal{K}$. Subsequently, nature commits to the portion of observed uncertain parameters $\boldsymbol{w} \circ \overline{\boldsymbol{\xi}}$ and to a choice $\boldsymbol{\xi}^k$, $k \in \mathcal{K}$, associated with each candidate policy $\boldsymbol{y}^k$. Finally, the decision-maker chooses one of the candidate policies.

In what follows, we provide insights into the theoretical and computational properties of the $K$-adaptability counterpart to two-stage robust problems with DDID and with binary recourse.

**Remark 2.** We note that the results in Section 3 generalize fully to cases where the objective and constraint functions are continuous (not necessarily linear) in $\boldsymbol{x}$, $\boldsymbol{y}$, and $\boldsymbol{\xi}$. Moreover, all of the ideas in our paper generalize to the case where the technology and recourse matrices, $\boldsymbol{T}$ and $\boldsymbol{W}$, depend on $\boldsymbol{\xi}$. We do not discuss these cases in detail so as to minimize notational overhead.

### 4.2. *K*-Adaptability for Problems with Objective Uncertainty

In this section, we focus our attention on the case where uncertain parameters only appear in the objective of Problem ($\mathcal{P}$) and where the recourse decisions are binary, being expressible as

$$\begin{aligned} \text{minimize} \quad & \max_{\overline{\boldsymbol{\xi}} \in \Xi} \min_{\boldsymbol{y} \in \mathcal{Y}} \left\{ \max_{\boldsymbol{\xi} \in \Xi(\boldsymbol{w}, \overline{\boldsymbol{\xi}})} \boldsymbol{\xi}^\top \boldsymbol{C}\boldsymbol{x} + \boldsymbol{\xi}^\top \boldsymbol{D}\boldsymbol{w} + \boldsymbol{\xi}^\top \boldsymbol{Q}\boldsymbol{y} : \boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y} \le \boldsymbol{h} \right\} \qquad (\mathcal{PO}) \\ \text{subject to} \quad & \boldsymbol{x} \in \mathcal{X},\ \boldsymbol{w} \in \mathcal{W}, \end{aligned}$$

where $\boldsymbol{h} \in \mathbb{R}^L$, $\mathcal{Y} \subseteq \{0,1\}^{N_y}$. We study the $K$-adaptability counterpart of Problem ($\mathcal{PO}$) given by

$$\begin{aligned} \text{minimize} \quad & \max_{\overline{\boldsymbol{\xi}} \in \Xi} \min_{k \in \mathcal{K}} \left\{ \max_{\boldsymbol{\xi} \in \Xi(\boldsymbol{w}, \overline{\boldsymbol{\xi}})} \boldsymbol{\xi}^\top \boldsymbol{C}\boldsymbol{x} + \boldsymbol{\xi}^\top \boldsymbol{D}\boldsymbol{w} + \boldsymbol{\xi}^\top \boldsymbol{Q}\boldsymbol{y}^k : \boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \le \boldsymbol{h} \right\} \qquad (\mathcal{PO}_K) \\ \text{subject to} \quad & \boldsymbol{x} \in \mathcal{X},\ \boldsymbol{w} \in \mathcal{W},\ \boldsymbol{y}^k \in \mathcal{Y},\ k \in \mathcal{K}. \end{aligned}$$

Applying Lemma 1, we are able to write Problem ($\mathcal{PO}_K$) equivalently as

$$\begin{aligned} \text{minimize} \quad & \max_{\{\boldsymbol{\xi}^k\}_{k\in\mathcal{K}} \in \Xi^K(\boldsymbol{w})} \min_{k \in \mathcal{K}} \{(\boldsymbol{\xi}^k)^\top \boldsymbol{C}\boldsymbol{x} + (\boldsymbol{\xi}^k)^\top \boldsymbol{D}\boldsymbol{w} + (\boldsymbol{\xi}^k)^\top \boldsymbol{Q}\boldsymbol{y}^k : \boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \le \boldsymbol{h}\} \\ \text{subject to} \quad & \boldsymbol{x} \in \mathcal{X},\ \boldsymbol{w} \in \mathcal{W},\ \boldsymbol{y}^k \in \mathcal{Y},\ k \in \mathcal{K}, \end{aligned} \qquad (7)$$

where $\Xi^K(\boldsymbol{w})$ is defined as in Lemma 1. In the absence of uncertainty in the constraints, the constraints in the $K$-adaptability problem can be moved to the first stage, as summarized by the following observation.

**Observation 1.** The $K$-adaptability counterpart of the two-stage robust optimization problem with decision-dependent information discovery, Problem ($\mathcal{PO}_K$), is equivalent to

$$\begin{aligned} \text{minimize} \quad & \max_{\{\boldsymbol{\xi}^k\}_{k\in\mathcal{K}} \in \Xi^K(\boldsymbol{w})} \min_{k \in \mathcal{K}} \{(\boldsymbol{\xi}^k)^\top \boldsymbol{C}\boldsymbol{x} + (\boldsymbol{\xi}^k)^\top \boldsymbol{D}\boldsymbol{w} + (\boldsymbol{\xi}^k)^\top \boldsymbol{Q}\boldsymbol{y}^k\} \\ \text{subject to} \quad & \boldsymbol{x} \in \mathcal{X},\ \boldsymbol{w} \in \mathcal{W},\ \mathbf{y}^k \in \mathcal{Y},\ k \in \mathcal{K} \\ & \boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \le \boldsymbol{h} \quad \forall k \in \mathcal{K}, \end{aligned} \qquad (8)$$

where $\Xi^K(\boldsymbol{w})$ is as defined in Equation (6).

Note that for all $\boldsymbol{w} \in \mathcal{W}$, the set $\Xi^K(\boldsymbol{w})$ is nonempty and bounded. Thus, $(\boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}^k\}_{k\in\mathcal{K}}) \in \mathcal{X} \times \mathcal{W} \times \mathcal{Y}^K$ is feasible in Problem (8) if $\boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \le \boldsymbol{h}$ for all $k \in \mathcal{K}$, whereas to be feasible in Problem (7) (and accordingly in Problem ($\mathcal{PO}_K$)), it need only satisfy $\boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \le \boldsymbol{h}$ for some $k \in \mathcal{K}$. Thus, a triplet $(\boldsymbol{x}, \boldsymbol{w}, \boldsymbol{y}^k)$ feasible in (7) (and thus in ($\mathcal{PO}_K$)) need not be feasible in Problem (8). However, the proof of Observation 1, provides a way to construct a feasible solution for Problem (8) from a feasible solution to Problem (7) that achieves the same optimal value.

Lemma 1 and Observation 1 are key to reformulating Problem ($\mathcal{PO}_K$) as a finite program. They also enable us to analyze the complexity of evaluating the objective function of the $K$-adaptability problem under a fixed decision. Indeed, from Problem (8), it can be seen that for any fixed choice $(\boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}^k\}_{k\in\mathcal{K}})$, the objective value of ($\mathcal{PO}_K$) can be evaluated by solving a linear program (LP) obtained by writing (8) in epigraph form. We formalize this result in the following.

**Observation 2.** For any fixed $K$ and decision $(\boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}^k\}_{k\in\mathcal{K}})$, the objective value of the $K$-adaptability problem ($\mathcal{PO}_K$) can be evaluated in polynomial time in the size of the input.

In Observation 2, we showed that for any fixed $K$, $\boldsymbol{x}$, $\boldsymbol{w}$, and $\boldsymbol{y}^k$, the objective function in Problem ($\mathcal{PO}_K$) can be evaluated by means of a polynomially sized LP. By dualizing this LP, we can obtain an equivalent reformulation of Problem ($\mathcal{PO}_K$) in the form of a bilinear problem.

**Theorem 2.** *Problem ($\mathcal{PO}_K$) is equivalent to the bilinear problem*

$$\begin{aligned} \text{minimize} \quad & \boldsymbol{b}^\top \boldsymbol{\beta} + \sum_{k\in\mathcal{K}} \boldsymbol{b}^\top \boldsymbol{\beta}^k \\ \text{subject to} \quad & \boldsymbol{x} \in \mathcal{X},\ \boldsymbol{w} \in \mathcal{W},\ \boldsymbol{y}^k \in \mathcal{Y},\ k \in \mathcal{K} \\ & \boldsymbol{\alpha} \in \mathbb{R}^K_+,\ \boldsymbol{\beta} \in \mathbb{R}^R_+,\ \boldsymbol{\beta}^k \in \mathbb{R}^R_+,\ \boldsymbol{\gamma}^k \in \mathbb{R}^{N_\xi},\ k \in \mathcal{K} \\ & \mathbf{e}^\top \boldsymbol{\alpha} = 1 \\ & \boldsymbol{A}^\top \boldsymbol{\beta}^k + \boldsymbol{w} \circ \boldsymbol{\gamma}^k = \boldsymbol{\alpha}_k (\boldsymbol{C}\boldsymbol{x} + \boldsymbol{D}\boldsymbol{w} + \boldsymbol{Q}\boldsymbol{y}^k) \quad \forall k \in \mathcal{K} \\ & \boldsymbol{A}^\top \boldsymbol{\beta} = \sum_{k\in\mathcal{K}} \boldsymbol{w} \circ \boldsymbol{\gamma}^k \\ & \boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \le \boldsymbol{h} \quad \forall k \in \mathcal{K}. \end{aligned} \qquad (9)$$

Although Problem ($\mathcal{PO}_K$) is generally nonconvex (bilinear), there exist several techniques in the literature for solving such problems exactly. In fact, this is an extremely active area of research, see for example, Tsoukalas and Mitsos (2014) and Gupte et al. (2017). Moreover, problems of the form ($\mathcal{PO}_K$) can now be solved with state-of-the-art off-the-shelf solvers like Gurobi. Indeed, Gurobi recently released its 9th version that can tackle nonconvex quadratic programs.[2] If $\mathcal{X} \subseteq \{0,1\}^{N_x}$ and $\mathcal{Y} \subseteq \{0,1\}^{N_y}$, the bilinear terms in the formulation above can be linearized using standard techniques and we can obtain an equivalent reformulation of Problem ($\mathcal{PO}_K$) in the form of an MBLP.

**Corollary 1.** *Suppose $\mathcal{X} \subseteq \{0,1\}^{N_x}$ and $\mathcal{Y} \subseteq \{0,1\}^{N_y}$. Then, Problem ($\mathcal{PO}_K$) is equivalent an MBLP involving a suitably chosen "big-M" constant.*

We emphasize that the size of the MBLP in Corollary 1 is polynomial in the size of the input data for the $K$-adaptability problem ($\mathcal{PO}_K$). Note that, contrary to Hanasusanto et al. (2015), to reformulate Problem ($\mathcal{PO}_K$) as an MBLP, we require that $\mathcal{X} \subseteq \{0,1\}^{N_x}$. This is to ensure that we are able to linearize the bilinear terms involving the $\boldsymbol{x}$ variables that arise from the dualization step. We note that formulation (9) and its equivalent MBLP can be augmented with symmetry breaking constraints to speed-up solution, see Section EC.5.1 for details.

**Remark 3.** Most MBLP solvers[3] allow reformulating the bilinear terms without the use of "big-$M$" constants, which are known to suffer from numerical instability. These include, for example, so-called `SOS` or `IfThen` constraints.

**Observation 3.** Suppose that we are only in the presence of exogenous uncertainty, that is, $\boldsymbol{w} = \mathbf{e}$, $\boldsymbol{D} = \mathbf{0}$, and $\boldsymbol{V} = \mathbf{0}$. Then, Problem (11) reduces to the MBLP formulation of the $K$-adaptability problem with only exogenous uncertainty from Hanasusanto et al. (2015).

A generalization of our model and solution approach in this section to the multistage case with objective uncertainty is provided in Electronic Companion EC.1.

### 4.3. *K*-Adaptability for Problems with Constraint Uncertainty

The starting point of our analysis is the reformulation of Problem ($\mathcal{P}_K$) as the min-max-min problem (5). Unfortunately, this problem is generally hard as testified by the following theorem.

**Theorem 3.** *Evaluating the objective of Problem (5) if $K$ is not fixed is strongly NP-hard.*

We reformulate Problem (5) equivalently by shifting the second-stage constraints $\boldsymbol{Tx} + \boldsymbol{Vw} + \boldsymbol{Wy}^k \le \boldsymbol{H\xi}^k$ from the objective function to the definition of the uncertainty set. We thus replace $\Xi^K(\boldsymbol{w})$ with a family of uncertainty sets parameterized by a vector $\boldsymbol{\ell}$.

**Proposition 1.** *The $K$-adaptability problem with decision-dependent information discovery, Problem (5), is equivalent to*

$$
\begin{array}{ll}
\text{minimize} & \displaystyle \max_{\boldsymbol{\ell}\in\mathcal{L}} \max_{\{\boldsymbol{\xi}^k\}_{k\in\mathcal{K}}\in\Xi^K(\boldsymbol{w},\boldsymbol{\ell})} \min_{\substack{k\in\mathcal{K}:\\ \ell_k=0}} \{(\boldsymbol{\xi}^k)^\top \boldsymbol{C}\boldsymbol{x} + (\boldsymbol{\xi}^k)^\top \boldsymbol{D}\boldsymbol{w} + (\boldsymbol{\xi}^k)^\top \boldsymbol{Q}\boldsymbol{y}^k\} \\
\text{subject to} & \boldsymbol{x}\in\mathcal{X},\ \boldsymbol{w}\in\mathcal{W},\ \boldsymbol{y}^k\in\mathcal{Y},\ k\in\mathcal{K},
\end{array} \tag{10}
$$

*where $\mathcal{L} := \{0,\dots,L\}^K$, $L$ is the number of second-stage constraints in Problem ($\mathcal{P}$), and the uncertainty sets $\Xi^K(\boldsymbol{w},\boldsymbol{\ell})$, $\boldsymbol{\ell}\in\mathcal{L}$, are defined as*

$$
\Xi^K(\boldsymbol{w},\boldsymbol{\ell}) := \left\{ \begin{array}{ll}
\{\boldsymbol{\xi}^k\}_{k\in\mathcal{K}} \in \Xi^K : & \\
\boldsymbol{w}\circ\boldsymbol{\xi}^k = \boldsymbol{w}\circ\bar{\boldsymbol{\xi}} & \forall k\in\mathcal{K} \text{ for some } \bar{\boldsymbol{\xi}}\in\Xi \\
\boldsymbol{Tx}+\boldsymbol{Vw}+\boldsymbol{Wy}^k \le \boldsymbol{H\xi}^k & \forall k\in\mathcal{K}: \ell_k = 0 \\
[\boldsymbol{Tx}+\boldsymbol{Vw}+\boldsymbol{Wy}^k]_{\ell_k} > [\boldsymbol{H\xi}^k]_{\ell_k} & \forall k\in\mathcal{K}: \ell_k \neq 0
\end{array} \right\},
$$

*where, for convenience, we have suppressed the dependence of $\Xi^K(\boldsymbol{w},\boldsymbol{\ell})$ on $\boldsymbol{x}$ and $\boldsymbol{y}^k$, $k\in\mathcal{K}$.*

The elements of vector $\boldsymbol{\ell}\in\mathcal{L}$ in Proposition 1 encode which second-stage policies are feasible for the parameter realizations $\{\boldsymbol{\xi}^k\}_{k\in\mathcal{K}} \in \Xi^K(\boldsymbol{w},\boldsymbol{\ell})$. Indeed, recall that $\boldsymbol{\xi}^k$ can be viewed as the recourse action that nature will take if the decision-maker acts according to $\boldsymbol{y}^k$ in response to seeing $\bar{\boldsymbol{\xi}}$. Thus, policy $\boldsymbol{y}^k$ is feasible in Problem (5) (and thus in Problem ($\mathcal{P}_K$)) if $\ell_k = 0$. On the other hand, policy $\boldsymbol{y}^k$ violates the $\ell_k$-th constraint in Problem (5) if $\ell_k \neq 0$. Thus, if $\ell_k \neq 0$, this implies that the $\ell_k$-th constraint in ($\mathcal{P}_K$) is violated for some $\boldsymbol{\xi}\in\Xi(\boldsymbol{w},\bar{\boldsymbol{\xi}})$ and therefore $\boldsymbol{y}^k$ is not feasible in ($\mathcal{P}_K$). Note that, in contrast to the case with exogenous uncertainty discussed by Hanasusanto et al. (2016), $\ell_k = 0$ if and only if policy $\boldsymbol{y}^k$ is *robustly* feasible in ($\mathcal{P}_K$).

Having brought Problem ($\mathcal{P}_K$) to the form (10), it now presents a similar structure to a problem with objective uncertainty (see Section 4.2) with the caveats that the problem involves multiple uncertainty sets that are also open. Next, we employ closed inner approximations $\Xi^K_\epsilon(\boldsymbol{w},\boldsymbol{\ell})$ of the sets $\Xi^K(\boldsymbol{w},\boldsymbol{\ell})$ that are parameterized by a scalar $\epsilon > 0$:

$$
\begin{array}{ll}
\text{minimize} & \displaystyle \max_{\boldsymbol{\ell}\in\mathcal{L}} \max_{\{\boldsymbol{\xi}^k\}_{k\in\mathcal{K}}\in\Xi^K_\epsilon(\boldsymbol{w},\boldsymbol{\ell})} \min_{\substack{k\in\mathcal{K}:\\ \ell_k=0}} \{(\boldsymbol{\xi}^k)^\top \boldsymbol{C}\boldsymbol{x} + (\boldsymbol{\xi}^k)^\top \boldsymbol{D}\boldsymbol{w} + (\boldsymbol{\xi}^k)^\top \boldsymbol{Q}\boldsymbol{y}^k\} \\
\text{subject to} & \boldsymbol{x}\in\mathcal{X},\ \boldsymbol{w}\in\mathcal{W},\ \mathbf{y}^k\in\mathcal{Y},\ k\in\mathcal{K},
\end{array} \tag{$10_\epsilon$}
$$

where the uncertainty sets $\Xi_\epsilon^K(\boldsymbol{w},\boldsymbol{\ell})$ are defined as

$$\Xi_\epsilon^K(\boldsymbol{w},\boldsymbol{\ell}) := \left\{ \begin{array}{ll} \{\boldsymbol{\xi}^k\}_{k\in\mathcal{K}} \in \Xi^K : & \\ \boldsymbol{w}\circ\boldsymbol{\xi}^k = \boldsymbol{w}\circ\bar{\boldsymbol{\xi}} & \forall k\in\mathcal{K} \text{ for some } \bar{\boldsymbol{\xi}}\in\Xi \\ \boldsymbol{T}\boldsymbol{x}+\boldsymbol{V}\boldsymbol{w}+\boldsymbol{W}\boldsymbol{y}^k \le \boldsymbol{H}\boldsymbol{\xi}^k & \forall k\in\mathcal{K}: \ell_k = 0 \\ [\boldsymbol{T}\boldsymbol{x}+\boldsymbol{V}\boldsymbol{w}+\boldsymbol{W}\boldsymbol{y}^k]_{\ell_k} \ge [\boldsymbol{H}\boldsymbol{\xi}^k]_{\ell_k} + \epsilon & \forall k\in\mathcal{K}: \ell_k \neq 0 \end{array} \right\}.$$

Using this definition, we next reformulate the approximate Problem ($10_\epsilon$) equivalently as an MBLP.

**Theorem 4.** *The approximate problem ($10_\epsilon$) is equivalent to the mixed binary bilinear program*

$$\begin{array}{ll}
\min & \tau \\
\text{s.t.} & \tau\in\mathbb{R},\ \boldsymbol{x}\in\mathcal{X},\ \boldsymbol{w}\in\mathcal{W},\ \boldsymbol{y}^k\in\mathcal{Y},\ k\in\mathcal{K} \\
& \boldsymbol{\alpha}(\boldsymbol{\ell})\in\mathbb{R}_+^R,\ \boldsymbol{\alpha}^k(\boldsymbol{\ell})\in\mathbb{R}_+^R,\ k\in\mathcal{K},\ \gamma(\boldsymbol{\ell})\in\mathbb{R}_+^K, \\
& \qquad \eta^k(\boldsymbol{\ell})\in\mathbb{R}^{N_\xi},\ k\in\mathcal{K},\ \boldsymbol{\ell}\in\mathcal{L} \\
& \left.\begin{array}{l}
\lambda(\boldsymbol{\ell})\in\Lambda_K(\boldsymbol{\ell}),\ \boldsymbol{\beta}^k(\boldsymbol{\ell})\in\mathbb{R}_+^L,\ k\in\mathcal{K}, \\
\boldsymbol{A}^\top\boldsymbol{\alpha}(\boldsymbol{\ell}) = \sum_{k\in\mathcal{K}} \boldsymbol{w}\circ\boldsymbol{\eta}^k(\boldsymbol{\ell}) \\
\boldsymbol{A}^\top\boldsymbol{\alpha}^k(\boldsymbol{\ell}) - \boldsymbol{H}^\top\boldsymbol{\beta}^k(\boldsymbol{\ell}) + \boldsymbol{w}\circ\eta^k(\boldsymbol{\ell}) \\
\quad = \lambda_k(\boldsymbol{\ell})\left[\boldsymbol{C}\boldsymbol{x}+\boldsymbol{D}\,\boldsymbol{w}+\boldsymbol{Q}\,\boldsymbol{y}^k\right] \quad \forall k\in\mathcal{K}: \ell_k = 0 \\
\boldsymbol{A}^\top\boldsymbol{\alpha}^k(\boldsymbol{\ell}) + [\boldsymbol{H}]_{\ell_k}\gamma_k(\boldsymbol{\ell}) + \boldsymbol{w}\circ\eta^k(\boldsymbol{\ell}) \\
\quad = \lambda_k(\boldsymbol{\ell})\left[\boldsymbol{C}\boldsymbol{x}+\boldsymbol{D}\,\boldsymbol{w}+\boldsymbol{Q}\,\boldsymbol{y}^k\right] \quad \forall k\in\mathcal{K}: \ell_k \neq 0 \\
\tau \ge \boldsymbol{b}^\top\left(\boldsymbol{\alpha}(\boldsymbol{\ell}) + \sum_{k\in\mathcal{K}}\boldsymbol{\alpha}^k(\boldsymbol{\ell})\right) - \sum_{\substack{k\in\mathcal{K}:\\ \ell_k=0}} (\boldsymbol{T}\boldsymbol{x}+\boldsymbol{V}\boldsymbol{w}+\boldsymbol{W}\boldsymbol{y}^k)^\top \\
\quad \boldsymbol{\beta}^k(\boldsymbol{\ell}) + \sum_{\substack{k\in\mathcal{K}:\\ \ell_k\neq 0}} ([\boldsymbol{T}\boldsymbol{x}+\boldsymbol{V}\boldsymbol{w}+\boldsymbol{W}\boldsymbol{y}^k]_{\ell_k} - \epsilon)\gamma_k(\boldsymbol{\ell})
\end{array}\right\} \forall \boldsymbol{\ell}\in\partial\mathcal{L} \\
& \left.\begin{array}{l}
\boldsymbol{A}^\top\boldsymbol{\alpha}(\boldsymbol{\ell}) = \sum_{k\in\mathcal{K}} \boldsymbol{w}\circ\eta^k(\boldsymbol{\ell}) \\
\boldsymbol{A}^\top\boldsymbol{\alpha}^k(\boldsymbol{\ell}) + [\boldsymbol{H}]_{\ell_k}\gamma_k(\boldsymbol{\ell}) + \boldsymbol{w}\circ\eta^k(\boldsymbol{\ell}) = \boldsymbol{0} \quad \forall k\in\mathcal{K} \\
\boldsymbol{b}^\top\left(\boldsymbol{\alpha}(\boldsymbol{\ell}) + \sum_{k\in\mathcal{K}}\boldsymbol{\alpha}^k(\boldsymbol{\ell})\right) \\
\quad + \sum_{k\in\mathcal{K}} ([\boldsymbol{T}\boldsymbol{x}+\boldsymbol{V}\boldsymbol{w}+\boldsymbol{W}\boldsymbol{y}^k]_{\ell_k} - \epsilon)\gamma_k(\boldsymbol{\ell}) \le -1
\end{array}\right\} \forall \boldsymbol{\ell}\in\mathcal{L}_+,
\end{array} \tag{11}$$

*where* $\Lambda_K(\boldsymbol{\ell}) := \{\boldsymbol{\lambda}\in\mathbb{R}_+^K : \mathbf{e}^\top\boldsymbol{\lambda} = 1,\ \boldsymbol{\lambda}_k = 0\ \forall k\in\mathcal{K}: \ell_k \neq 0\}$, $\partial\mathcal{L} := \{\boldsymbol{\ell}\in\mathcal{L} : \boldsymbol{\ell} \ngtr \boldsymbol{0}\}$ *and* $\mathcal{L}_+ := \{\boldsymbol{\ell}\in\mathcal{L} : \boldsymbol{\ell} > \boldsymbol{0}\}$ *denote the sets for which the decision* $(\boldsymbol{x},\boldsymbol{w},\{\boldsymbol{y}_k\}_{k\in\mathcal{K}})$ *satisfies or violates the second-stage constraints in Problem (10), respectively.*

If $\mathcal{X}\subseteq\{0,1\}^{N_x}$ and $\mathcal{Y}\subseteq\{0,1\}^{N_y}$, then for $\epsilon$ sufficiently small, $\cup_{\boldsymbol{\ell}\in\mathcal{L}}\Xi_\epsilon^K(\boldsymbol{w},\boldsymbol{\ell})$ is nonempty for all $(\boldsymbol{x},\boldsymbol{w},\{\boldsymbol{y}^k\}_{k\in\mathcal{K}})$ feasible in Problem ($10_\epsilon$), implying that Problem ($10_\epsilon$) is bounded below. Therefore, for $\epsilon$ sufficiently small, its equivalent Problem (11) is also bounded below and thus admits an equivalent reformulation as an MBLP involving a suitably chosen "big-$M$" constant. Similar to the robust counterpart resulting from the decision rule approximation proposed in Vayanos et al. (2011), Problem (11) presents a number of constraints and decision variables that is *exponential* in the approximation parameter, in this case $K$. Relative to the prepartitioning approach from Vayanos et al. (2011), our method does however present a number of distinct advantages. First, the trade-off between approximation quality and computational tractability is controlled using a *single* design parameter; in contrast, in the prepartitioning approach, the number of design parameters equals the number of observable uncertain parameters. Second, as we increase $K$, the quality of the approximation improves in our case, whereas increasing the number of breakpoints along a given direction does not necessarily yield to improvements in the prepartitioning approach. Finally, to identify breakpoint configurations resulting in low optimality gap, a large number of optimization problems need to be solved.

**Remark 4.** Theorem 4 directly generalizes to instances of Problem ($\mathcal{P}_K$) where the technology and recourse matrices $\boldsymbol{T}$, $\boldsymbol{V}$, and $\boldsymbol{W}$ depend on $\boldsymbol{\xi}$. Indeed, it suffices to absorb the coefficients of any uncertain terms in $\boldsymbol{T}$, $\boldsymbol{V}$, and $\boldsymbol{W}$ in the right-hand side matrix $\boldsymbol{H}$.

**Observation 4.** Suppose that we are only in the presence of exogenous uncertainty, that is, $\boldsymbol{w} = \mathbf{e}$, $\boldsymbol{D} = \boldsymbol{0}$, and $\boldsymbol{V} = \boldsymbol{0}$. Then, Problem (11) reduces to the MBLP formulation of the $K$-adaptability problem with constraint uncertainty and with only exogenous uncertain parameters from Hanasusanto et al. (2015). In particular, in the case of constraint uncertainty, Hanasusanto et al. (2015) also require that the first stage variables $\boldsymbol{x}$ be binary.

## 5. The Case of Piecewise Linear Convex Objective

In this section, we investigate two-stage robust optimization problems with DDID and objective uncertainty where the objective function is given as the maximum of finitely many linear functions.

### 5.1. Problem Formulation

A piecewise linear convex objective function can be written compactly as the maximum of finitely many linear functions of $\boldsymbol{\xi}$ and $(\boldsymbol{x},\boldsymbol{w},\boldsymbol{y})$, being expressible as

$$\max_{i\in\mathcal{I}}\ \boldsymbol{\xi}^\top\boldsymbol{C}^i\boldsymbol{x} + \boldsymbol{\xi}^\top\boldsymbol{D}^i\boldsymbol{w} + \boldsymbol{\xi}^\top\boldsymbol{Q}^i\boldsymbol{y}, \tag{12}$$

where $\boldsymbol{C}^i\in\mathbb{R}^{N_\xi\times N_x}$, $\boldsymbol{D}^i\in\mathbb{R}^{N_\xi\times N_\xi}$, and $\boldsymbol{Q}^i\in\mathbb{R}^{N_\xi\times N_y}$, $i\in\mathcal{I}$, $\mathcal{I}\subseteq\mathbb{N}$. A two-stage robust optimization problem with DDID, objective function given by (12), and objective uncertainty is then expressible as

$$\begin{array}{ll}
\min & \max_{\bar{\boldsymbol{\xi}}\in\Xi}\ \min_{\boldsymbol{y}\in\mathcal{Y}}\ \max_{\boldsymbol{\xi}\in\Xi(\boldsymbol{w},\bar{\boldsymbol{\xi}})} \left\{\max_{i\in\mathcal{I}}\ \boldsymbol{\xi}^\top\boldsymbol{C}^i\boldsymbol{x} + \boldsymbol{\xi}^\top\boldsymbol{D}^i\boldsymbol{w} + \boldsymbol{\xi}^\top\boldsymbol{Q}^i\boldsymbol{y}\right\} \\
\text{s.t.} & \boldsymbol{x}\in\mathcal{X},\ \boldsymbol{w}\in\mathcal{W}.
\end{array} \tag{$\mathcal{PO}^{\text{PWL}}$}$$

Note that, as in Section 4.2, our framework remains applicable in the presence of joint deterministic constraints on the first and second stage variables. We omit these to minimize notational overhead.

### 5.2. *K*-Adaptability Approximation & MBLP Reformulation

The $K$-adaptability counterpart of Problem ($\mathcal{PO}^{\text{PWL}}$) reads

$$\begin{aligned} &\min \max_{\bar{\boldsymbol{\xi}}\in\Xi} \min_{k\in\mathcal{K}} \max_{\boldsymbol{\xi}\in\Xi(\boldsymbol{w},\bar{\boldsymbol{\xi}})} \left\{ \max_{i\in\mathcal{I}} \boldsymbol{\xi}^\top \boldsymbol{C}^i \boldsymbol{x} + \boldsymbol{\xi}^\top \boldsymbol{D}^i \boldsymbol{w} + \boldsymbol{\xi}^\top \boldsymbol{Q}^i \boldsymbol{y}^k \right\} \\ &\text{s.t.} \quad \boldsymbol{x}\in\mathcal{X}, \boldsymbol{w}\in\mathcal{W}, \boldsymbol{y}^k\in\mathcal{Y}, k\in\mathcal{K}. \end{aligned} \qquad (\mathcal{PO}_K^{\text{PWL}})$$

We begin this reformulation by the following lemma, which parallels Lemma 1, and shows that we can exchange the order of the inner min and max in formulation ($\mathcal{PO}_K^{\text{PWL}}$), by indexing $\boldsymbol{\xi}$ by $k$.

**Lemma 2.** *The K-adaptability counterpart of Problem ($\mathcal{PO}_K^{\text{PWL}}$) is equivalent to*

$$\begin{aligned} \text{minimize} \quad & \max_{\{\boldsymbol{\xi}^k\}_{k\in\mathcal{K}}\in\Xi^K(\boldsymbol{w})} \min_{k\in\mathcal{K}} \left\{ \max_{i\in\mathcal{I}} (\boldsymbol{\xi}^k)^\top \boldsymbol{C}^i \boldsymbol{x} + (\boldsymbol{\xi}^k)^\top \boldsymbol{D}^i \boldsymbol{w} + (\boldsymbol{\xi}^k)^\top \boldsymbol{Q}^i \boldsymbol{y}^k \right\} \\ \text{subject to} \quad & \boldsymbol{x}\in\mathcal{X}, \boldsymbol{w}\in\mathcal{W}, \boldsymbol{y}^k\in\mathcal{Y}, k\in\mathcal{K}. \end{aligned} \qquad (13)$$

Next, by leveraging Lemma 2, we are able to reformulate Problem (13) exactly as an MBLP.

**Theorem 5.** *Problem ($\mathcal{PO}_K^{\text{PWL}}$) is equivalent to the bilinear program*

$$\begin{aligned} \text{minimize} \quad & \tau \\ \text{subject to} \quad & \tau\in\mathbb{R}, \boldsymbol{x}\in\mathcal{X}, \boldsymbol{w}\in\mathcal{W}, \boldsymbol{y}^k\in\mathcal{Y}, k\in\mathcal{K} \\ & \boldsymbol{\alpha}^{\boldsymbol{i}}\in\mathbb{R}^K_+, \boldsymbol{\beta}^{\boldsymbol{i}}\in\mathbb{R}^R_+, \boldsymbol{\beta}^{\boldsymbol{i},k}\in\mathbb{R}^R_+, \boldsymbol{\gamma}^{\boldsymbol{i},k}\in\mathbb{R}^{N_\xi}, \quad \forall k\in\mathcal{K}, \boldsymbol{i}\in\mathcal{I}^K \\ & \left.\begin{aligned} & \tau \geq \boldsymbol{b}^\top\boldsymbol{\beta}^{\boldsymbol{i}} + \sum_{k\in\mathcal{K}} \boldsymbol{b}^\top \boldsymbol{\beta}^{\boldsymbol{i},k} \\ & \mathbf{e}^\top \boldsymbol{\alpha}^{\boldsymbol{i}} = 1 \\ & \boldsymbol{A}^\top\boldsymbol{\beta}^{\boldsymbol{i},k} + \boldsymbol{w}\circ\boldsymbol{\gamma}^{\boldsymbol{i},k} \\ & \quad = \boldsymbol{\alpha}^{\boldsymbol{i}}_k(\boldsymbol{C}^{\boldsymbol{i}_k}\boldsymbol{x} + \boldsymbol{D}^{\boldsymbol{i}_k}\boldsymbol{w} + \boldsymbol{Q}^{\boldsymbol{i}_k}\boldsymbol{y}^k) \quad \forall k\in\mathcal{K} \\ & \boldsymbol{A}^\top\boldsymbol{\beta}^{\boldsymbol{i}} = \sum_{k\in\mathcal{K}} \boldsymbol{w}\circ\boldsymbol{\gamma}^{\boldsymbol{i},k} \end{aligned}\right\} \forall \boldsymbol{i}\in\mathcal{I}^K, \end{aligned} \qquad (14)$$

*which can be written as an MBLP, provided* $\mathcal{X}\subseteq\{0,1\}^{N_x}$.

Albeit Problem (14) is an MBLP, it presents an exponential number of decision variables and constraints making it difficult to solve directly using off-the-shelf solvers even when $K$ is only moderately large ($K \gtrsim 4$). In the remainder of this section, we exploit the specific structure of Problem ($\mathcal{PO}^{\text{PWL}}$) to solve its $K$-adaptability counterpart *exactly* by reformulating it as an MBLP that presents an attractive structure amenable to decomposition techniques.

### 5.3. "Column-and-Constraint Generation" Algorithm

Column-and-constraint generation techniques are a popular approach for addressing problems that possess an exponential number of decision variables and constraints while presenting a decomposable structure, see for example, Fischetti and Vigo (1997), Löbel (1998), Valério De Carvalho (1999), Mamer and McBride (2000), Feillet et al. (2010), Sadykov and Vanderbeck (2011), Zeng and Zhao (2013), Muter et al. (2013), and Muter et al. (2018). We propose a new column-and-constraint generation algorithm to solve the $K$-adaptability counterpart ($\mathcal{PO}_K^{\text{PWL}}$) based on its reformulation (14). The key idea is to decompose the problem into a relaxed master problem and a series of subproblems indexed by $\boldsymbol{i}\in\mathcal{I}^K$. The master problem initially only involves the first stage constraints and a *single auxiliary MBLP* is used to iteratively identify indices $\boldsymbol{i}\in\mathcal{I}^K$ for which the solution to the relaxed master problem becomes infeasible when plugged into subproblem $\boldsymbol{i}$. Constraints associated with infeasible subproblems are added to the master problem and the procedure continues until convergence. We detail this procedure in Electronic Companion EC.3 where we also show that certain classes of two-stage robust optimization problems that seek to minimize the "worst-case absolute regret" criterion can be written in the form ($\mathcal{PO}^{\text{PWL}}$). In Section 7, we leverage the column and constraint generation algorithm and this observation to solve an active preference elicitation problem that seeks to recommend kidney allocation policies with least possible worst-case regret.

## 6. Computational Studies on Stylized Instances

We investigate the performance of our approach on a variety of two-stage robust optimization problems with decision-dependent information discovery (for computational results on multistage robust optimization with decision-dependent information discovery, see Electronic Companion EC.2). We solve these problems with our proposed methods discussed in Sections 4.2, 4.3, and 5.2. To speed-up computation, for the two-stage problems, we employ a conservative greedy heuristic that uses the solution to problems with smaller $K$ to solve problems with larger $K$ more efficiently, see Electronic Companion EC.5.2. This strategy enables us to solve many random instances of problems with large approximation parameters $K$ (up to $K = 10$). In all our experiments, we compare our method to the state-of-the-art prepartitioning approach from Vayanos et al. (2011) using the ROC++ platform, see Vayanos et al. (2023). All of our experiments are performed on the

High Performance Computing Cluster of our university. Each job is allotted 64GB of RAM, 16 cores, and a 2.6GHz Xeon processor. All optimization problems are solved using Gurobi v9.0.1. We allow a total time limit of 7,200 seconds to solve each instance cumulatively across all values of $K$ for the $K$-adaptability problem and across all breakpoint configurations for the prepartitioning approach. Unless indicated otherwise, we set $M = 1{,}000$ for the $K$-adaptability approach. We found that this value yielded good performance in most cases.

### 6.1. Two-Stage Robust Best Box Selection (Objective Uncertainty)

The first problem we study is a robust variant of the best box selection problem, see for example, Gupta et al. (2016, 2017) for results on the stochastic version. In this problem, an agent must select one out of $N$ boxes, indexed in the set $\mathcal{N} := \{1, \ldots, N\}$, each of which contains a prize. The value $\boldsymbol{\xi}_i \in \mathbb{R}$ of the prize in each box $i \in \mathcal{N}$ is unknown and will only be revealed if the box is opened. Opening box $i \in \mathcal{N}$ incurs a cost $\boldsymbol{c}_i$. In the first stage, the agent can decide whether to open each box $i \in \mathcal{N}$ which we indicate with the decision variables $w_i \in \{0,1\}$. Thus, $w_i = 1$ if and only if $\boldsymbol{\xi}_i$ is observed between the first and second decision stages. The total budget available to open boxes is $B$. In the second stage, the agent can choose one of the opened boxes to keep, which we indicate with the decision variable $\boldsymbol{y}_i \in \{0,1\}$, $i \in \mathcal{N}$, earning its prize. We assume that the value of box $i \in \mathcal{N}$ is expressible as $\boldsymbol{\xi}_i = (1 + \boldsymbol{\Phi}_i^\top \boldsymbol{\zeta}/2)\boldsymbol{\xi}_i^0$, where $\boldsymbol{\xi}_i^0$ corresponds to the nominal value of the prize of box $i$, $\boldsymbol{\zeta} \in [-1,1]^L$ are $L$ risk factors, and $\boldsymbol{\Phi}_i \in \mathbb{R}^L$ collects the factor loadings associated with the value of box $i$. The goal of the agent is to select the boxes to open (first stage decisions) and the box to keep (second stage decision) to maximize the worst-case value of the box kept. The Best Box Selection problem has numerous applications, for example in house purchasing or in candidate interviewing, see for example, Singla (2018). With the notation above, the problem can be expressed as a two-stage robust optimization problem with decision-dependent information discovery of the form ($\mathcal{PO}$) as

$$\max_{w \in \{0,1\}^N} \min_{\overline{\xi} \in \Xi} \max_{y \in \{0,1\}^N} \left\{ \min_{\xi \in \Xi(w, \overline{\xi})} \boldsymbol{\xi}^\top \boldsymbol{y} : \mathbf{e}^\top \boldsymbol{y} = 1, \boldsymbol{c}^\top \boldsymbol{w} \le B, \boldsymbol{y} \le \boldsymbol{w} \right\}, \tag{15}$$

where $\Xi := \{\boldsymbol{\xi} \in \mathbb{R}^N : \exists \boldsymbol{\zeta} \in [-1,1]^L : \boldsymbol{\xi}_i = (1 + \boldsymbol{\Phi}_i^\top \boldsymbol{\zeta}/2)\boldsymbol{\xi}_i^0,\ i = 1, \ldots, N\}$.

We evaluate the performance of our approach on 100 randomly generated instances of Problem (15) with $L = 4$ risk factors: 20 instances for each $N \in \{10, 20, 30, 40, 50\}$. In these instances, $\boldsymbol{c}$ is drawn uniformly at random from the box $[0, 10]^N$, we let $\boldsymbol{\xi}_i^0 = \boldsymbol{c}/5$, and $B = \mathbf{e}^\top \boldsymbol{c}/2$. The matrix $\boldsymbol{\Phi}$ is sampled uniformly at random from the box $[-1,1]^{N \times L}$. Our computational results across those instances are summarized in Table 1. From the table, we observe that with the proposed $K$-adaptability approach, all instances (even those involving $N = 50$ boxes) solved to optimality with an average solver time no greater than 2 seconds across all problem sizes. In contrast, the average solver time of the prepartitioning approach exceeded 190 seconds for $N = 20$ boxes and equaled 6,713 seconds for $N = 50$ boxes, with only 95.7% of the problems associated with all breakpoint configurations solving within the allotted time on average. In addition, the quality of the best solution identified by the proposed $K$-adaptability solution consistently outperformed that of the best prepartitioning solution. For example, an average improvement of over

**Table 1.** Summary of Computational Results on the Best Box Selection Problem for Various Choices of $N$ over 20 Randomly Generated Instances of Each Size

| | Adapt. | $N = 10$, $L = 4$ | $N = 20$, $L = 4$ | $N = 30$, $L = 4$ | $N = 40$, $L = 4$ | $N = 50$, $L = 4$ |
|---|---|---|---|---|---|---|
| $K$-adaptability | $K = 1$ | 100%/0.0%/0s | 100%/0.0%/0s | 100%/0.0%/0s | 100%/0.0%/0s | 100%/0.0%/0s |
| | $K = 2$ | 100%/98.3%/0s | 100%/83.6%/0s | 100%/81.0%/0s | 100%/61.2%/0s | 100%/59.1%/0s |
| | $K = 3$ | 100%/136.9%/0s | 100%/127.5%/0s | 100%/110.4%/0s | 100%/81.1%/0s | 100%/94.0%/0s |
| | $K = 4$ | 100%/165.3%/0s | 100%/150.0%/0s | 100%/127.3%/1s | 100%/91.0%/1s | 100%/113.6%/0s |
| | $K = 5$ | 100%/167.9%/0s | 100%/157.8%/1s | 100%/137.1%/2s | 100%/98.4%/1s | 100%/117.0%/1s |
| | $K = 6$ | 100%/170.4%/0s | 100%/162.2%/1s | 100%/144.6%/3s | 100%/102.4%/2s | 100%/122.6%/1s |
| | $K = 7$ | 100%/170.4%/0s | 100%/162.6%/1s | 100%/147.5%/3s | 100%/109.2%/2s | 100%/126.1%/1s |
| | $K = 8$ | 100%/170.4%/0s | 100%/162.8%/1s | 100%/148.2%/3s | 100%/111.2%/2s | 100%/127.3%/2s |
| | $K = 9$ | 100%/170.4%/0s | 100%/162.8%/1s | 100%/148.3%/3s | 100%/111.6%/2s | 100%/128.0%/2s |
| | $K = 10$ | 100%/170.4%/0s | 100%/162.8%/1s | 100%/148.3%/3s | 100%/111.6%/2s | 100%/128.4%/2s |
| Prepartitioning | ≤10 subsets | 100%/155.8%/ 17s/7.5 | 100%/142.8%/ 190s/8.1 | 100%/125.6%/ 978s/8.1 | 100%/96.7%/ 2,785s/8.6 | 95.7%/115.5%/ 6,713s/8.2 |

*Notes.* In the $K$-adaptability part of the table, each entry corresponds to: percentage of instances solved within the time limit/average improvement in the objective value of the $K$-adaptable solution over the static solution/average solution time across all instances. In the prepartitioning part of the table, each entry corresponds to: average percentage of breakpoint configurations that solved within the time limit out of all configurations with cardinality at most 10/average improvement in the objective value of the best prepartitioning solution found within the time limit relative to that of the static solution/average cumulative solver time/average cardinality of the best solution found within the time limit.

148% over the static solution was exhibited for the $K$-adaptability method for $N = 30$, while the prepartitioning solution only resulted in a 125.6% improvement. Finally, the smallest value of $K$ needed to achieve the same average performance as the best solution in the prepartitioning method was always smaller than the average number of subsets needed to obtain that solution in the prepartitioning approach, resulting in more easy to use/implement solutions for our proposed method. For example, for $N = 30$, 8.1 subsets are needed on average to obtain an improvement of 125.6% for the prepartitioning approach, whereas $K = 4$ candidate strategies suffice for our proposed $K$-adaptability approach to yield an improvement of 127.3%.

### 6.2. Preference Elicitation with Real-Valued Recommendations (Real Decisions)

The second problem we consider is a robust active preference elicitation problem where user preferences can be elicited by asking them "how much" they like any particular item and where real-valued quantities of multiple items can be recommended after preferences are elicited, see Vayanos et al. (2022) for a variant where pairwise comparison queries are used instead.

The building blocks of our framework are candidate items which we index in the set $\mathcal{I} := \{1, \ldots, I\}$. We let $\boldsymbol{\phi}^i \in \mathbb{R}^J$ be the feature vector of item $i \in \mathcal{I}$. We assume that user preferences are cardinal and model them by means of a linear utility function. Specifically, we assume that the utility of item $i$ is given by $u(\boldsymbol{\phi}^i) = \boldsymbol{u}^\top \boldsymbol{\phi}^i + \tilde{\epsilon}_i$, where $\{\tilde{\epsilon}_i\}_{i \in \mathcal{I}}$ are independent identically distributed and $\boldsymbol{u}$ is a vector of (unknown) utility function coefficients supported in the set $\mathcal{U} \subseteq [-1, 1]^J$. These assumptions are standard in the literature, see for example, Bertsimas and O'Hair (2013) and Boutilier et al. (2004). Before making recommendations, the system has the opportunity to make $Q$ queries to the user. Each query is based on one of the candidate items: if query $i \in \mathcal{I}$ is chosen, the user is asked "On a scale from 0 to 1, where 1 is the most anyone could like an item and 0 is the least anyone could like an item, how much do you like policy $i$?" We denote by $\boldsymbol{\xi}_i \in [0, 1]$ the answer to query $i$. After the answers to these queries are observed, the system can select $N$ out of the $I$ items to recommend and the quantity $\boldsymbol{y}_i \in [0, 1]$, $i \in \mathcal{I}$, of those items to recommend. The goal of the recommender system is to select $Q$ queries the answers to which will enable the system to recommend a set of items in quantities resulting in greatest possible worst-case utility.

To formulate the preference elicitation problem mathematically we let $\boldsymbol{w}_i$, $i \in \mathcal{I}$, denote the decision to pose query $i$, that is, to observe $\boldsymbol{\xi}_i$ before making a recommendation. Thus, $\mathcal{W} := \{\boldsymbol{w} \in \{0, 1\}^I : \mathbf{e}^\top \boldsymbol{w} = Q\}$. The set of possible realizations of $\boldsymbol{\xi}$ is given by

$$\Xi := \left\{ \boldsymbol{\xi} \in [0, 1]^I : \exists \boldsymbol{u} \in [-1, 1]^J, \boldsymbol{\epsilon} \in \mathcal{E} \text{ such that } \boldsymbol{\xi}_i = \frac{\boldsymbol{u}^\top \boldsymbol{\phi}^i + \max_{j \in \mathcal{I}} \|\boldsymbol{\phi}^j\|_1}{2 \max_{j \in \mathcal{I}} \|\boldsymbol{\phi}^j\|_1} + \epsilon_i \quad \forall i \in \mathcal{I} \right\},$$

where the normalization of $\boldsymbol{u}^\top \boldsymbol{\phi}^i$ ensures that $\boldsymbol{\xi}_i$ has the correct interpretation and, in the spirit of modern robust optimization, see for example, Lorca and Sun (2016), we assume that $\boldsymbol{\epsilon}\psi$is valued in the set $\mathcal{E} := \{\boldsymbol{\epsilon} \in \mathbb{R}^I : \sum_{i=1}^I |\epsilon_i| \leq \Gamma\}$, where $\Gamma$ is a user-specified *budget of uncertainty* parameter. Once the answers to the queries are observed, the recommender system may select the quantity of each item $i \in \mathcal{I}$ to recommend which we encode with decisions $\boldsymbol{y}_i \in [0, 1]$. We let $z_i \in \{0, 1\}$ indicate if item $i$ is recommended and require that the quantity of items recommended equals 1. Thus,

$$\mathcal{Y} := \{\boldsymbol{y} \in [0, 1]^I : \exists \boldsymbol{z} \in \{0, 1\}^I \text{ such that } \mathbf{e}^\top \boldsymbol{z} = N,\ \boldsymbol{y} \leq \boldsymbol{z},\ \mathbf{e}^\top \boldsymbol{y} = 1\}.$$

With this notation, the preference elicitation problem is expressible as

$$\underset{\boldsymbol{w} \in \mathcal{W}}{\text{maximize}} \ \min_{\overline{\boldsymbol{\xi}} \in \Xi} \ \max_{\boldsymbol{y} \in \mathcal{Y}} \ \min_{\boldsymbol{\xi} \in \Xi(\boldsymbol{w}, \overline{\boldsymbol{\xi}})} \boldsymbol{\xi}^\top \boldsymbol{y}. \qquad (\mathcal{WCU}^{\text{PE}})$$

A conservative solution to Problem ($\mathcal{WCU}^{\text{PE}}$) can be obtained using the $K$-adaptability approximation scheme discussed in Section 4.2, by solving the bilinear reformulation (9).

We evaluate the performance of our approach on 80 randomly generated instances of Problem ($\mathcal{WCU}^{\text{PE}}$): 20 instances for each $(Q, N, \Gamma) \in \{(1, 2, 0.1), (3, 3, 0.3), (6, 4, 0.6), (9, 5, 0.9)\}$. In these instances, $I = 30$, $J = 15$, and $\boldsymbol{\phi}^i$, $i \in \mathcal{I}$, are drawn uniformly at random from the box $[-1, 1]^J$. Our computational results across these instances are summarized in Table 2. From the table, we observe that on average the optimal value of our proposed $K$-adaptability method (across all $K$) is greater than that of the best optimal value of the prepartitioning method (across all breakpoint configurations). For example, for the $(Q, N, \Gamma) = (9, 5, 0.9)$ setting, $K$-adaptability yields an average improvement in optimal value of 53.3% relative to the static solution, whereas prepartitioning only results in an average improvement of 10.1% on average in the best case. In addition, the solutions obtained by the $K$-adaptability approach in the same time needed to solve for all breakpoint configurations (or to reach the time limit) in the prepartitioning approach are of far better quality. For example, the prepartitioning approach always reached the 7,200 seconds time limit for instances of size $(Q, N, \Gamma) = (3, 3, 0.3)$ with an associated average improvement in optimal value of 12.6%. In contrast,

**Table 2.** Summary of Computational Results on the Preference Elicitation Problem with Real-Valued Recommendations for Various Choices of $Q$, $N$, and $\Gamma$ over 20 Randomly Generated Instances for Each Setting

| | Adapt. | $Q=1$, $N=2$, $\Gamma=0.1$ | $Q=3$, $N=3$, $\Gamma=0.3$ | $Q=6$, $N=4$, $\Gamma=0.6$ | $Q=9$, $N=5$, $\Gamma=0.9$ |
|---|---|---|---|---|---|
| $K$-adaptability | $K=1$ | 100%/0.0%/2s | 100%/0.0%/4s | 100%/0.0%/8s | 100%/0.0%/26s |
| | $K=2$ | 100%/14.0%/7s | 100%/28.2%/17s | 100%/35.7%/26s | 100%/41.7%/51s |
| | $K=3$ | 100%/15.2%/19s | 100%/29.0%/89s | 100%/42.6%/418s | 100%/52.2%/238s |
| | $K=4$ | 100%/16.1%/39s | 100%/29.4%/238s | 100%/43.5%/1,689s | 90%/53.0%/2,772s |
| | $K=5$ | 100%/16.8%/68s | 100%/29.4%/509s | 80%/43.9%/3,572s | 55%/53.2%/4,751s |
| | $K=6$ | 100%/16.9%/105s | 100%/29.4%/900s | 65%/44.1%/5,088s | 40%/53.2%/5,917s |
| | $K=7$ | 100%/17.3%/152s | 100%/29.6%/1,434s | 25%/44.2%/6,362s | 15%/53.2%/6,609s |
| | $K=8$ | 100%/17.3%/203s | 100%/29.8%/2,151s | 15%/44.2%/6,776s | 15%/53.3%/6,904s |
| | $K=9$ | 100%/17.3%/262s | 95%/30.0%/3,033s | 5%/44.2%/6,957s | 5%/53.3%/7,108s |
| | $K=10$ | 100%/17.5%/347s | 90%/30.2%/4,155s | 5%/44.2%/7,024s | 5%/53.3%/7,185s |
| Prepartitioning | ≤10 subsets | 3.5%/16.2%/7,200s/8.2 | 3.0%/12.6%/7,200s/9.7 | 0.1%/13.8%/7,200s/7.4 | 0.0%/10.1%/7,200s/6.8 |

*Note.* The row names and table entries have the same interpretation as in Table 1.

within just 17 seconds on average, the $K$-adaptability approach results in an improvement of 28.2% in optimal value on average over the same instances. Finally, the average value of $K$ needed to achieve a solution of quality comparable to that of the best prepartitioning approach is a lot smaller than the number of subsets needed in prepartitioning, implying that $K$-adaptability has more attractive usability properties. For example, for $(Q,N,\Gamma)=(6,4,0.6)$, 7.4 subsets are needed by prepartitioning to yield a 13.8% improvement in optimal value whereas $K=2$ is sufficient for our method to yield an improvement of 35.7%.

## 6.3. Robust R&D Project Portfolio Optimization (Constraint Uncertainty)

The third problem we investigate is a robust variant of the R&D project portfolio optimization problem, see for example, Solak et al. (2010) for a solution approach on the stochastic version. In this problem, an R&D firm has a pipeline of $N$ candidate projects indexed in the set $\mathcal{N}:=\{1,\ldots,N\}$ that it can invest in. The return $\boldsymbol{\xi}_i^{\mathrm{r}}$ of each project $i\in\mathcal{N}$ is uncertain and will only be revealed if the firm chooses to undertake the project. The firm can decide to undertake each project $i\in\mathcal{N}$ in year one, indicated by decision $w_i^{\mathrm{r}}\in\{0,1\}$, in the following year, indicated by decision $y_i\in\{0,1\}$, or not at all. Thus, $w_i^{\mathrm{r}}=1$ if and only if $\boldsymbol{\xi}_i^{\mathrm{r}}$ is observed between the first and second years. If the firm chooses to undertake the investment in the second year, it will only realize a known fraction $\theta\in(0,1]$ of the return. Undertaking project $i$ incurs an unknown cost $\boldsymbol{\xi}_i^{\mathrm{c}}$ that will only be revealed if the firm chooses to undertake the project. The total budget available to invest in projects across the two years is $B$. We assume that the return and cost of project $i\in\mathcal{N}$ are expressible as $\boldsymbol{\xi}_i^{\mathrm{r}}=(1+\boldsymbol{\Phi}_i^\top\boldsymbol{\zeta}/2)\boldsymbol{\xi}_i^{\mathrm{r},0}$ and $\boldsymbol{\xi}_i^{\mathrm{c}}=(1+\boldsymbol{\Psi}_i^\top\boldsymbol{\zeta}/2)\boldsymbol{\xi}_i^{\mathrm{c},0}$, where $\boldsymbol{\xi}_i^{\mathrm{r},0}$ and $\boldsymbol{\xi}_i^{\mathrm{c},0}$ corresponds to the nominal return and cost for project $i$, respectively, $\boldsymbol{\zeta}\in[-1,1]^L$ are $L$ risk factors, and the vectors $\boldsymbol{\Phi}_i\in\mathbb{R}^L$ and $\boldsymbol{\Psi}_i\in\mathbb{R}^L$ collect the factor loadings for the return and cost of project $i$, respectively.

With this notation, the R&D project portfolio optimization problem is expressible as a two-stage robust optimization problem with decision-dependent information discovery of the form ($\mathcal{P}$) as

$$\begin{aligned}\text{maximize}\quad & \min_{\bar{\boldsymbol{\xi}}\in\Xi}\ \max_{\boldsymbol{y}\in\{0,1\}^N}\left\{\min_{\boldsymbol{\xi}\in\Xi(\boldsymbol{w},\bar{\boldsymbol{\xi}})}(\boldsymbol{w}^{\mathrm{r}}+\theta\boldsymbol{y})^\top\boldsymbol{\xi}^{\mathrm{r}}:\right.\\ &\left.\qquad(\boldsymbol{w}^{\mathrm{r}}+\boldsymbol{y})^\top\boldsymbol{\xi}^{\mathrm{c}}\le B,\ \boldsymbol{w}^{\mathrm{r}}+\boldsymbol{y}\le\mathbf{e}\right\}\\ \text{subject to}\quad & \boldsymbol{w}=(\boldsymbol{w}^{\mathrm{r}},\boldsymbol{w}^{\mathrm{r}}),\ \boldsymbol{w}^{\mathrm{r}}\in\{0,1\}^N,\end{aligned} \tag{16}$$

where

$$\Xi:=\{(\boldsymbol{\xi}^{\mathrm{r}},\boldsymbol{\xi}^{\mathrm{c}})\in\mathbb{R}^{2N}:\exists\boldsymbol{\zeta}\in[-1,1]^L:\boldsymbol{\xi}_i^{\mathrm{r}}=(1+\boldsymbol{\Phi}_i^\top\boldsymbol{\zeta}/2)\boldsymbol{\xi}_i^{\mathrm{r},0},\ \boldsymbol{\xi}_i^{\mathrm{c}}=(1+\boldsymbol{\Psi}_i^\top\boldsymbol{\zeta}/2)\boldsymbol{\xi}_i^{\mathrm{c},0},\ i=1,\ldots,N\}.$$

In this problem, we set $M=100$ and $\epsilon=10^{-4}$. We evaluate the performance of our approach on 100 randomly generated instances of Problem (16): 20 instances for each $(N,L)\in\{(5,3),(10,5),(15,8),(20,10),(25,13)\}$. In these instances, $\theta=0.8$, $\boldsymbol{\xi}^{\mathrm{c},0}$ is drawn uniformly at random from the box $[0,10]^N$, and we let $\boldsymbol{\xi}^{\mathrm{r},0}=\boldsymbol{\xi}^{\mathrm{c},0}/5$ and $B=\mathbf{e}^\top\boldsymbol{\xi}^{\mathrm{c},0}/2$. The elements of $\boldsymbol{\Phi}$ and $\boldsymbol{\Psi}$ are uniformly distributed in the interval $[-1,1]$. We remove any instance where $K=1$ returns an optimal value of zero in the corresponding $K$-adaptability problem so that we can quantify the percentage improvement relative to this static solution. Our computational results across these instances are summarized in Table 3. From the table, we observe that on average the optimal value of our proposed $K$-adaptability method (across all $K$) is greater than that of the best optimal value of the prepartitioning method (across all breakpoint configurations). For example, for the $(N,L)=(10,5)$ setting, $K$-adaptability yields an average improvement in optimal value of 88.1% relative to the static solution, whereas prepartitioning only results in an average improvement of 39.3% in the best case. In addition, the solutions obtained by

**Table 3.** Summary of Computational Results on the R&D Project Portfolio Optimization Problem for Various Choices of $N$ and $L$ over 20 Randomly Generated Instances of Each Size

| | Adapt. | $N=5$, $L=3$ | $N=10$, $L=5$ | $N=15$, $L=8$ | $N=20$, $L=10$ | $N=25$, $L=13$ |
|---|---|---|---|---|---|---|
| $K$-adaptability | $K=1$ | 100%/0.0%/0s | 100%/0.0%/0s | 100%/0.0%/2s | 100%/0.0%/15s | 100%/0.0%/33s |
| | $K=2$ | 100%/71.9%/0s | 100%/40.7%/2s | 100%/40.7%/16s | 100%/29.9%/166s | 95%/44.2%/1,744s |
| | $K=3$ | 100%/114.4%/1s | 100%/60.4%/7s | 100%/55.2%/61s | 100%/43.6%/966s | 65%/60.8%/5,018s |
| | $K=4$ | 100%/119.8%/2s | 100%/67.9%/23s | 100%/62.6%/273s | 80%/49.7%/2,876s | 10%/68.6%/6,945s |
| | $K=5$ | 100%/125.6%/5s | 100%/73.7%/73s | 100%/69.1%/1,436s | 30%/53.4%/6,479s | 0%/69.4%/7,200s |
| | $K=6$ | 100%/125.7%/8s | 100%/78.7%/263s | 75%/72.4%/5,151s | 0%/54.3%/7,200s | 0%/69.4%/7,200s |
| | $K=7$ | 100%/131.4%/18s | 100%/83.7%/1,200s | 5%/74.1%/7,100s | 0%/54.3%/7,200s | 0%/69.4%/7,200s |
| | $K=8$ | 100%/131.5%/60s | 65%/87.0%/4,777s | 0%/74.1%/7,200s | 0%/54.3%/7,200s | 0%/69.4%/7,200s |
| | $K=9$ | 100%/137.1%/276s | 0%/88.1%/7,200s | 0%/74.1%/7,200s | 0%/54.3%/7,200s | 0%/69.4%/7,200s |
| | $K=10$ | 100%/137.3%/1,747s | 0%/88.1%/7,200s | 0%/74.1%/7,200s | 0%/54.3%/7,200s | 0%/69.4%/7,200s |
| Prepartitioning | ≤10 subsets | 100%/72.3%/ 16s/7.1 | 100%/39.3%/ 853s/8.4 | 92.4%/30.5%/ 6,675s/9.1 | 21.0%/21.2%/ 7,200s/8.8 | 3.9%/17.3%/ 7,200s/8.5 |

*Note.* The row names and table entries have the same interpretation as in Table 1.

the $K$-adaptability approach in the same time needed to solve for all breakpoint configurations (or to reach the time limit) in the prepartitioning approach are of far better quality. For example, the prepartitioning approach needed 6,675 seconds on average to solve instances of size $(N, L) = (15, 8)$ with an associated average improvement in optimal value of 30.5%. In contrast, within just 16 seconds on average, the $K$-adaptability approach results in an improvement of 40.7% in optimal value on average over the same instances. Finally, and similar to our results on the best box problem in Section 6.1, the $K$ adaptability approach has more attractive usability properties than prepartitioning. For example, for $(N, L) = (20, 10)$, 8.8 subsets are needed by prepartitioning to yield a 21.2% improvement in optimal value whereas $K = 2$ is sufficient for our method to yield an improvement of 29.9%.

# 7. Preference Elicitation to Improve the US Kidney Allocation System

In this section, we evaluate our approach on a preference elicitation and recommendation problem that explicitly captures the endogenous nature of the elicitation process.

## 7.1. Motivation & Problem Formulation (Piecewise Linear Convex Objective)

The motivation for our study is one of the central problems faced by policymakers at the OPTN/UNOS who must periodically make changes to the policy for prioritizing patients on the kidney transplant waiting list for scarce deceased donor kidneys. To tackle this problem, a Kidney Transplantation Committee (KTC) is appointed at the OPTN that examines the outcomes of numerous candidate policies simulated using the Kidney-Pancreas Simulated Allocation Model (KPSAM), a simulator developed by the SRTR, see KPSAM (2015). The KTC examines the outcomes of the allocation policy alternatives along several dimensions (measures) of fairness and efficiency (e.g., number of recipients by age group, number of deaths by gender) before ultimately committing to one of the alternatives. This process was for example followed in the latest big policy change, see for example, Wolfe et al. (2009). Since selecting one alternative (policy) over many others is a challenging task, in particular when the dimension of each alternative is large, see for example, Toubia et al. (2003, 2004, 2007) and Boutilier et al. (2004), we propose a preference elicitation and recommendation framework for identifying a preferred policy using a moderate number of strategically chosen queries.

We formulate this problem as a variant of the active preference elicitation problem from Section 6.2 where a single item can be recommended and where we select queries that minimize worst-case regret of the recommendation. Items indexed in the set $\mathcal{I}$ correspond to policies where the feature vector $\boldsymbol{\phi}^i \in \mathbb{R}^J$ of policy $i \in \mathcal{I}$ collects various measures of fairness and efficiency of the policy. The problem is expressible mathematically as

$$\underset{\boldsymbol{w}\in\mathcal{W}}{\text{minimize}} \ \max_{\overline{\boldsymbol{\xi}}\in\Xi} \ \min_{\boldsymbol{y}\in\mathcal{Y}} \ \max_{\boldsymbol{\xi}\in\Xi(\boldsymbol{w},\overline{\boldsymbol{\xi}})} \left\{ \max_{i\in\mathcal{I}} \ \boldsymbol{\xi}_i - \boldsymbol{\xi}^\top \boldsymbol{y} \right\}, \qquad (\mathcal{WCR}^{\text{PE}})$$

where $\mathcal{W}$ and $\Xi$ are as in Section 6.2 and where $\mathcal{Y} := \{\boldsymbol{y} \in \{0,1\}^I : \mathbf{e}^\top \boldsymbol{y} = 1\}$. In this problem, the first part of the objective computes the utility of the best item to offer in hindsight, after the utilities $\boldsymbol{\xi}$ have been observed. The second part of the objective corresponds to the worst-case utility of the item recommended when only a portion of the uncertain parameters are observed, as dictated by the vector $\boldsymbol{w}$. Problem ($\mathcal{WCR}^{\text{PE}}$) can be solved approximately using the $K$-adaptability approximation scheme discussed in Section 5. Indeed, the regret in Problem ($\mathcal{WCR}^{\text{PE}}$) is given as the maximum of finitely many linear functions and Theorem 5 applies.

We note that in this case $|\mathcal{Y}| = I$. Thus, solving the $K$-adaptability counterpart of ($\mathcal{WCR}^{\mathrm{PE}}$) with $K = I$ recovers an optimal solution to the corresponding original problem.

## 7.2. Generating KAS Candidate Policies

This study used data from the Scientific Registry of Transplant Recipients (SRTR). The SRTR data system includes data on all donor, wait-listed candidates, and transplant recipients in the US, submitted by the members of the Organ Procurement and Transplantation Network (OPTN). The Health Resources and Services Administration (HRSA), U.S. Department of Health and Human Services provides oversight to the activities of the OPTN and SRTR contractors.

We generate the outcomes $\boldsymbol{\phi}^i$, $i \in \mathcal{I}$, of $I = 20$ candidate policies using the KPSAM simulator which we obtained from the SRTR using a modeling window from 01/01/2010 to 12/31/2010. The candidate policies we consider are linear scoring rules that use the patient dialysis time, the life years from transplant score, the Calculated Panel Reactive Antibodies and the age of the patient. For each policy, we record $J = 22$ outcomes, including the number of transplants overall, by age, by blood type, by race, and by gender, and the number of deaths by race and by gender. For details on the construction of the policies and for a list of outcomes, see Electronic Companion EC.4.

## 7.3. Numerical Results on KAS Candidate Policies

We evaluate the performance of our approach on the KAS policies data set from Section 7.2. Throughout our experiments, the $K$-adaptability counterpart of Problem ($\mathcal{WCR}^{\mathrm{PE}}$) is solved using the techniques described in Section 5. To speed-up computation, we also use a heuristic adapted from Subramanyam et al. (2020) and detailed in Section EC.5.2. The tolerance $\delta$ used in the column-and-constraint generation algorithm (see Section 5.3) is $10^{-5}$. We evaluate the *true* worst-case regret of any given solution $\boldsymbol{w}^\star$, which we denote by $r_{\mathrm{wc}}(\boldsymbol{w}^\star)$, as follows: we fix $\boldsymbol{w} = \boldsymbol{w}^\star$ in

**Figure 2.** (Color online) Optimality-Scalability Results for the Min-Max Regret Preference Elicitation Problem ($\mathcal{WCR}^{\mathrm{PE}}$) on the KAS Data

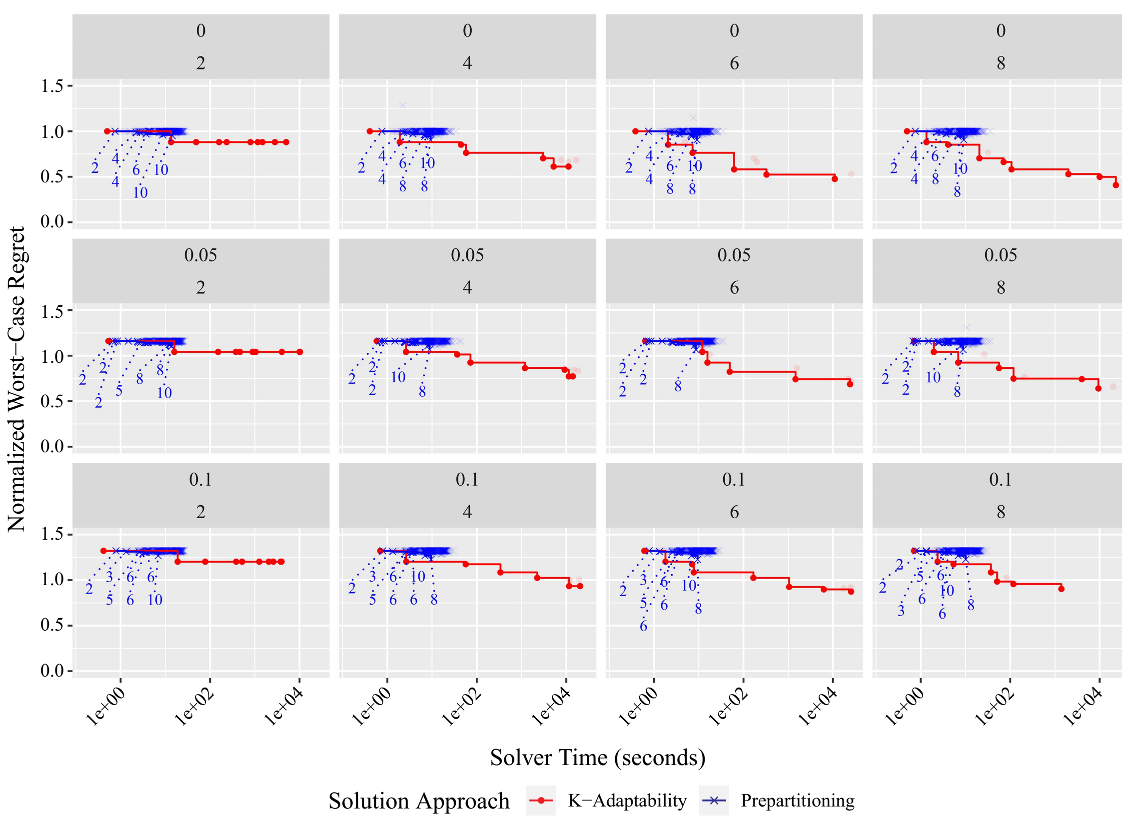


*Notes.* The numbers on each facet correspond to the values of $\Gamma$ (top number) and $Q$ (bottom number). Each dot corresponds to a different choice of $K \in \{1, \ldots, 10\}$ for the $K$-adaptability problem. Each cross corresponds to a different breakpoint configuration for the prepartitioning approach (we consider all breakpoint configurations drawn randomly from the set of all configurations with cardinality less than 10). The transparency of each dot and cross depends on whether it is on the efficient frontier of the configurations that resulted in the highest average optimal value for the given time budget. The numbers next to the efficient points indicate the degree of adaptability in the corresponding solution, i.e., the number of subsets in the prepartitioning approach.

**Figure 3.** (Color online) Results on the Performance of the $K$-Adaptability Approach Relative to Random Elicitation for the Min-Max Regret Preference Elicitation Problem ($\mathcal{WCR}^{\mathrm{PE}}$) on the KAS Data Set

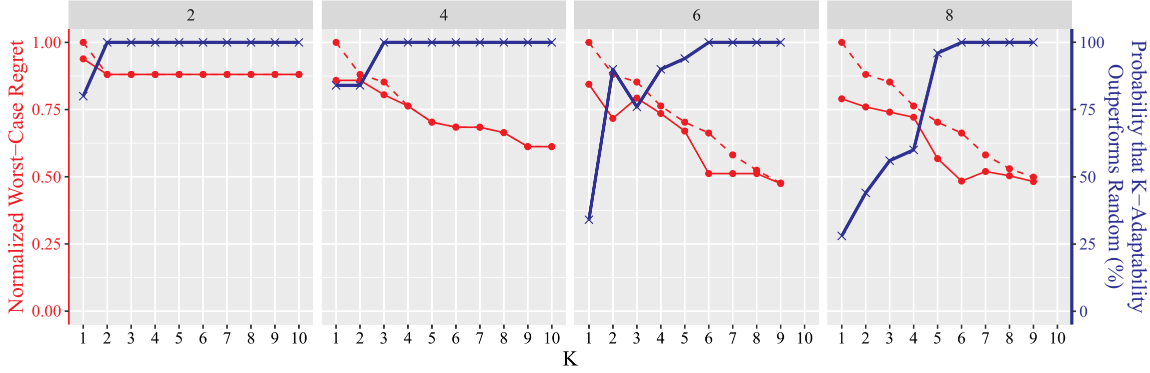


*Notes.* The number on each facet corresponds to the value of $Q$. The dashed thin line corresponds to the objective value of the $K$-adaptability problem. The thin solid line corresponds to $r_{\mathrm{wc}}(\boldsymbol{w}_K^\star)$ where $\boldsymbol{w}_K^\star$ is the optimal $K$-adaptable solution. The thick line represents the percentage of time that $r_{\mathrm{wc}}(\boldsymbol{w}_{\mathrm{r}})$ was lower than $r_{\mathrm{wc}}(\boldsymbol{w}_K^\star)$, where $\boldsymbol{w}_{\mathrm{r}}$ is a randomly drawn feasible solution.

Problem ($\mathcal{CCG}_{\mathrm{feas}}(\boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}^k\}_{k\in\mathcal{K}})$), where we set $K = I$ and employ all $I$ candidate policies $\{\boldsymbol{y}^k\}_{k\in\mathcal{K}}$ in the set $\mathcal{Y}$. As before, we use the ROC++ platform to solve the prepartitioning problem, see Vayanos et al. (2023). All of our experiments are performed using the same computing resources as in Section 6.

**7.3.1. Optimality-Scalability Trade-Off.** We evaluate the trade-off between computational complexity and scalability of our approach. We solve the min-max regret problems as $Q$ and $\Gamma$ are varied in the sets $\{2, 4, 6, 8\}$ and $\{0, 0.05, 0.1\}$, respectively. The results are summarized in Figure 2. From the figure it can be seen that the $K$-adaptability approach significantly outperforms the prepartitioning approach and static policies are very suboptimal. In fact, the prepartitioning approach performs comparably to static policies across all settings. On the other hand, with the $K$-adaptability approach, the normalized[4] worst-case regret drops to 0.40, 0.68, and 0.9 from 1, 1.16, and 1.32, for $\Gamma = 0, 0.05$, and 0.1, respectively (for $Q = 8$). This experiment shows the strength of the $K$-adaptability approach compared with the state of the art.

**7.3.2. Performance Relative to Random Elicitation.** We evaluate the benefits of computing near-optimal queries using the $K$-adaptability approximation approach relative to asking questions at random. We compare the *true* performance of a solution to the $K$-adaptability problem, $r_{\mathrm{wc}}(\boldsymbol{w}_K^\star)$, to that of 50 questions drawn uniformly at random from the set $\mathcal{W}$, $r_{\mathrm{wc}}(\boldsymbol{w}_{\mathrm{r}})$. The results are summarized on Figure 3. From the figure, we see that the probability that the $K$-adaptability solution outperforms random elicitation converges to 1 as $K$ grows. We observe that, for values of $K$ greater than 5, the $K$-adaptability solution outperforms random elicitation in over 90% of the cases.

## Disclaimer

The data reported here have been supplied by the Hennepin Healthcare Research Institute (HHRI) as the contractor for the Scientific Registry of Transplant Recipients (SRTR). The interpretation and reporting of these data are the responsibility of the author(s) and in no way should be seen as an official policy of or interpretation by the SRTR or the U.S. Government.

## Acknowledgments

The authors are grateful to Miss. Qing Jin for valuable discussions on implementation issues. The authors are also grateful to all referees whose comments helped substantially improve the quality of this paper.

## Endnotes

[1] See https://www.srtr.org, https://optn.transplant.hrsa.gov, and https://unos.org.

[2] See for example, https://www.gurobi.com/documentation/9.0/refman/nonconvex.html.

[3] See for example, https://www.ibm.com/analytics/cplex-optimizer and https://www.gurobi.com/.

[4] To aid with interpretability, we normalize regret such that the worst-case regret when no question is asked ($Q = 0$) and there is no error ($\Gamma = 0$) is 1 and the worst-case regret when all questions are asked is 0.

## References

Aghaei S, Azizi MJ, Vayanos P (2019) Learning optimal and fair decision trees for non-discriminative decision-making. *Proc. 33rd AAAI Conf. Artificial Intelligence* (AAAI, Washington, DC).

Aghaei S, Gómez A, Vayanos P (2024) Strong optimal classification trees. *Oper. Res.*, ePub ahead of print July 31, https://doi.org/10.1287/opre.2021.0034.

Assavapokee T, Realff MJ, Ammons JC (2008a) Min-max regret robust optimization approach on interval data uncertainty. *J. Optim. Theory Appl.* 137:297–316.

Assavapokee T, Realff MJ, Ammons JC, Hong IH (2008b) Scenario relaxation algorithm for finite scenario-based min-max regret and min-max relative regret robust optimization. *Comput. Oper. Res.* 35(6):2093–2102.

Ayoub J, Poss M (2016) Decomposition for adjustable robust linear optimization subject to uncertainty polytope. *Comput. Management Sci.* 13(2):219–239.

Basciftci B, Ahmed S, Shen S (2021) Distributionally robust facility location problem under decision-dependent stochastic demand. *Eur. J. Oper. Res.* 292(2):548–561.

Ben-Tal A, Nemirovski A (1998) Robust convex optimization. *Math. Oper. Res.* 23(4):769–805.

Ben-Tal A, Nemirovski A (1999) Robust solutions of uncertain linear programs. *Oper. Res. Lett.* 25(1):1–13.

Ben-Tal A, Nemirovski A (2000) Robust solutions of linear programming problems contaminated with uncertain data. *Math. Programming* 88:411–424.

Ben-Tal A, El Ghaoui L, Nemirovski A (2009) *Robust Optimization*, Princeton Series in Applied Mathematics (Princeton University Press, Princeton, NJ).

Ben-Tal A, Goryashko A, Guslitzer E, Nemirovski A (2004) Adjustable robust solutions of uncertain linear programs. *Math. Programming* 99(2):351–376.

Bertsimas D, Caramanis C (2010) Finite adaptability for linear optimization. *IEEE Trans. Automatic Control* 55(12):2751–2766.

Bertsimas D, Dunn J (2017) Optimal classification trees. *Machine Learn.* 106(7):1039–1082.

Bertsimas D, Dunning I (2016) Multistage robust mixed-integer optimization with adaptive partitions. *Oper. Res.* 64(4): 980–998.

Bertsimas D, Georghiou A (2015) Design of near optimal decision rules in multistage adaptive mixed-integer optimization. *Oper. Res.* 63(3):610–627.

Bertsimas D, Georghiou A (2018) Binary decision rules for multistage adaptive mixed-integer optimization. *Math. Programming* 167(2):395–433.

Bertsimas D, Goyal V (2012) On the power and limitations of affine policies in two-stage adaptive optimization problems. *Math. Programming* 134:491–531.

Bertsimas D, O'Hair A (2013) Learning preferences under noise and loss aversion: An optimization approach. *Oper. Res.* 61(5): 1190–1199.

Bertsimas D, Shtern S (2018) A scalable algorithm for two-stage adaptive linear optimization. Technical report, Operations Research Center, MIT, Cambridge, MA.

Bertsimas D, Sim M (2004) The price of robustness. *Oper. Res.* 52(1):35–53.

Bertsimas D, Vayanos P (2017) Data-driven learning in dynamic pricing using adaptive robust optimization. *Optimi. Online* (October 11), https://optimization-online.org/2014/10/4595/.

Bertsimas D, Iancu D, Parrilo P (2011) A hierarchy of near-optimal policies for multi-stage adaptive optimization. *IEEE Trans. Automatic Control* 56(12):2809–2824.

Bertsimas D, Pachamanova D, Sim M (2004) Robust linear optimization under general norms. *Oper. Res. Lett.* 32(6):510–516.

Bertsimas D, Delarue A, Jaillet P, Martin S (2019) The price of interpretability. Preprint, submitted July 8, https://arxiv.org/abs/1907.03419.

Boutilier C, Sandholm T, Shields R (2004) Eliciting bid taker non-price preferences in (combinatorial) auctions. *Proc. 19th Conf. Artificial Intelligence* (ACM, New York), 204–211.

Chassein A, Goerigk M, Kurtz J, Poss M (2019) Faster algorithms for min-max-min robustness for combinatorial problems with budgeted uncertainty. *Eur. J. Oper. Res.* 279(2):308–319.

Chen B, Wang J, Wang L, He Y, Wang Z (2014) Robust optimization for transmission expansion planning: Minimax cost vs. minimax regret. *IEEE Trans. Power Systems* 29(6):3069–3077.

Colvin M, Maravelias CT (2008) A stochastic programming approach for clinical trial planning in new drug development. *Comput. Chemical Engrg.* 32(11):2626–2642.

Colvin M, Maravelias CT (2010) Modeling methods and a branch and cut algorithm for pharmaceutical clinical trial planning using stochastic programming. *Eur. J. Oper. Res.* 203(1):205–215.

Feillet D, Gendreau M, Medaglia AL, Walteros JL (2010) A note on branch-and-cut-and-price. *Oper. Res. Lett.* 38(5):346–353.

Fischetti M, Vigo D (1997) A branch-and-cut algorithm for the resource-constrained minimum-weight arborescence problem. *Networks* 29(1):55–67.

Goel V, Grossman IE (2004) A stochastic programming approach to planning of offshore gas field developments under uncertainty in reserves. *Comput. Chemical Engrg.* 28(8):1409–1429.

Goel V, Grossman IE (2006) A class of stochastic programs with decision dependent uncertainty. *Math. Programming* 108(2): 355–394.

Goel V, Grossman IE, El-Bakry AA, Mulkay EL (2006) A novel branch and bound algorithm for optimal development of gas fields under uncertainty in reserves. *Comput. Chemical Engrg.* 30(6–7):1076–1092.

Gupta V, Grossmann IE (2011) Solution strategies for multistage stochastic programming with endogenous uncertainties. *Comput. Chemical Engrg.* 35(11):2235–2247.

Gupta A, Nagarajan V, Singla S (2016) Algorithms and adaptivity gaps for stochastic probing. *Proc. Annual ACM-SIAM Symposium Discrete Algorithms*, vol. 3 (ACM, New York).

Gupta A, Nagarajan V, Singla S (2017) Adaptivity gaps for stochastic probing: Submodular and XOS functions. *Proc. Annual ACM-SIAM Symposium Discrete Algorithms* (ACM, New York).

Gupte A, Ahmed S, Dey SS, Cheon MS (2017) Relaxations and discretizations for the pooling problem. *J. Global Optim.* 67:631–669.

Hanasusanto GA, Kuhn D, Wiesemann W (2015) K-adaptability in two-stage robust binary programming. *Oper. Res.* 63(4):877–891.

Hanasusanto GA, Kuhn D, Wiesemann W (2016) K-adaptability in two-stage distributionally robust binary programming. *Oper. Res. Lett.* 44(1):6–11.

Jiang R, Wang J, Zhang M, Guan Y (2013) Two-stage minimax regret robust unit commitment. *IEEE Trans. Power Systems* 28(3):2271–2282.

Jonsbråten TW (1998) Optimization models for petroleum field exploitation. PhD thesis, Norwegian School of Economics and Business Administration, Bergen, Norway.

Jonsbråten TW, Wets RB, Woodruff DL (1998) A class of stochastic programs with decision dependent random elements. *Ann. Oper. Res.* 82:83–106.

Koç A, Morton DP (2015) Prioritization via stochastic optimization. *Management Sci.* 61(3):586–603.

KPSAM (2015) Kidney-Pancreas Simulated Allocation Model User Guide version 2015. https://www.srtr.org/media/1202/kpsam.pdf.

Löbel A (1998) Vehicle scheduling in public transit and Lagrangean pricing. *Management Sci.* 44(12 Part 1):1637–1649.

Lorca A, Sun XA (2016) Multistage robust unit commitment with dynamic uncertainty sets and energy storage. *IEEE Trans. Power Systems* 32(3):1678–1688.

Luo F, Mehrotra S (2020) Distributionally robust optimization with decision dependent ambiguity sets. *Optim. Lett.* 14(8): 2565–2594.

Mamer JW, McBride RD (2000) Decomposition-based pricing procedure for large-scale linear programs: An application to the linear multicommodity flow problem. *Management Sci.* 46(5):693–709.

McCarthy SM, Laan CM, Wang K, Vayanos P, Sinha A, Tambe M (2018) The price of usability: Designing operationalizable strategies for security games. *Proc. Twenty-Seventh Internat. Joint Conf. Artificial Intelligence, IJCAI-18* (IJCAI, California), 454–460.

Muter I, Birbil SI, Bülbül K (2013) Simultaneous column-and-row generation for large-scale linear programs with column-dependent-rows. *Math. Programming* 142:47–82.

Muter I, Birbil I, Bülbül K (2018) Benders decomposition and column-and-row generation for solving large-scale linear programs with column-dependent-rows. *Eur. J. Oper. Res.* 264(1):29–45.

Ng TS (2013) Robust regret for uncertain linear programs with application to co-production models. *Eur. J. Oper. Res.* 227(3): 483–493.

Ning C, You F (2018) Adaptive robust optimization with minimax regret criterion: Multiobjective optimization framework and computational algorithm for planning and scheduling under uncertainty. *Comput. Chemical Engrg.* 108:425–447.

Nohadani O, Roy A (2017) Robust optimization with time-dependent uncertainty in radiation therapy. *IISE Trans. Healthcare Systems Engrg.* 7(2):81–92.

Nohadani O, Sharma K (2018) Optimization under decision-dependent uncertainty. *SIAM J. Optim.* 28(2):1773–1795.

Noyan N, Rudolf G, Lejeune M (2022) Distributionally robust optimization under a decision-dependent ambiguity set with applications to machine scheduling and humanitarian logistics. *INFORMS J. Comput.* 34(2):729–751.

Postek K, Den Hertog D (2016) Multistage adjustable robust mixed-integer optimization via iterative splitting of the uncertainty set. *INFORMS J. Comput.* 28(3):553–574.

Poursoltani M, Delage E (2019) Adjustable robust optimization reformulations of two-stage worst-case regret minimization problems. *Optim. Online* (July 22), https://optimization-online.org/2019/07/7303/.

Rahmattalabi A, Vayanos P, Fulginiti A, Rice E, Wilder B, Yadav A, Tambe M (2019) Exploring algorithmic fairness in robust graph covering problems. *Proc. 33rd Conf. Neural Inform. Processing Systems (NeurIPS)* (Curran Associates Inc., Red Hook, NY).

Sadykov R, Vanderbeck F (2011) Column generation for extended formulations. *Electronic Notes Discrete Math.* 37:357–362.

Shapiro A (2017) Interchangeability principle and dynamic equations in risk averse stochastic programming. *Oper. Res. Lett.* 45(4):377–381.

Singla S (2018) Combinatorial optimization under uncertainty: Probing and stopping-time algorithms. PhD thesis, Carnegie Mellon University, Pittsburgh.

Solak S, Clarke JP, Johnson EL, Barnes ER (2010) Optimization of R&D project portfolios under endogenous uncertainty. *Eur. J. Oper. Res.* 207(1):420–433.

Spacey SA, Wiesemann W, Kuhn D, Luk W (2012) Robust software partitioning with multiple instantiation. *INFORMS J. Comput.* 24(3):500–515.

Subramanyam A, Gounaris CE, Wiesemann W (2020) K-adaptability in two-stage mixed-integer robust optimization. *Math. Programming Comput.* 12:193–224.

Toubia O, Hauser J, Garcia R (2007) Probabilistic polyhedral methods for adaptive choice-based conjoint analysis: Theory and application. *Marketing Sci.* 26(5):596–610.

Toubia O, Hauser JR, Simester DI (2004) Polyhedral methods for adaptive choice-based conjoint analysis. *J. Marketing Res.* 41(1): 116–131.

Toubia O, Simester DI, Hauser JR, Dahan E (2003) Fast polyhedral adaptive conjoint estimation. *Marketing Sci.* 22(3):273–303.

Tsoukalas A, Mitsos A (2014) Multivariate McCormick relaxations. *J. Global Optim.* 59:633–662.

Valério De Carvalho JM (1999) Exact solution of bin-packing problems using column generation and branch-and-bound. *Ann. Oper. Res.* 86:629–659.

Vayanos P, Jin Q, Elissaios G (2023) ROC++: Robust optimization in C++. *INFORMS J. Comput.* 34(6):2873–2888.

Vayanos P, Kuhn D, Rustem B (2011) Decision rules for information discovery in multi-stage stochastic programming. *Proc. 50th IEEE Conf. Decision Control* (IEEE, Piscataway, NJ), 7368–7373.

Vayanos P, Kuhn D, Rustem B (2012) A constraint sampling approach for multi-stage robust optimization. *Automatica* 48(3): 459–471.

Vayanos P, Ye Y, McElfresh D, Dickerson J, Rice E (2022) Robust active preference elicitation. Preprint, submitted March 4, https://arxiv.org/abs/2003.01899.

Wolfe R, Leichtman A, McCullough K, Rodgers A (2009) Final analyses for data requests from the OPTN Kidney Transplantation Committee Meeting of August 24, 2009. Technical report, Arbor Research/University of Michigan, Ann Arbor, MI.

Xu G, Burer S (2018) A copositive approach for two-stage adjustable robust optimization with uncertain right-hand sides. *Comput. Optim. Appl.* 70(1):33–59.

Yu X, Shen S (2022) Multistage distributionally robust mixed-integer programming with decision-dependent moment-based ambiguity sets. *Math. Programming* 196(1–2):1025–1064.

Zeng B, Zhao L (2013) Solving two-stage robust optimization problems using a column-and-constraint generation method. *Oper. Res. Lett.* 41(5):457–461.

Zhang M (2011) Two-stage minimax regret robust uncapacitated lot-sizing problems with demand uncertainty. *Oper. Res. Lett.* 39(5): 342–345.

Zhang X, Kamgarpour M, Georghiou A, Goulart P, Lygeros J (2017) Robust optimal control with adjustable uncertainty sets. *Automatica* 75(Supplement C):249–259.

Zhen J, Den Hertog D, Sim M (2018) Adjustable robust optimization via Fourier-Motzkin elimination. *Oper. Res.* 66(4): 1086–1100.