Robust Optimization with Decision-Dependent Information Discovery

Phebe Vayanos
University of Southern California, Center for Artificial Intelligence in Society, phebe.vayanos@usc.edu

Angelos Georghiou
McGill, Desautels Faculty of Management, angelos.georghiou@mcgill.ca

Han Yu
University of Southern California, Center for Artificial Intelligence in Society, hyu376@usc.edu

Robust optimization is a popular paradigm for modeling and solving two- and multi-stage decision-making problems affected by uncertainty. Most approaches assume 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. 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. To fill this gap, we consider two- and multi-stage 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 mixed-binary linear program solvable with off-the-shelf solvers. We generalize our approach to the minimization of piecewise linear convex functions. We demonstrate effectiveness of our approach on synthetic and real data instances of the active preference elicitation problem used to recommend policies that meet the needs of policy-makers at the Los Angeles Homeless Services Authority.

Key words: robust optimization, endogenous uncertainty, decision-dependent information discovery, integer programming, two-stage problems, preference elicitation, worst-case regret.

1. Introduction

1.1. Background & Motivation

Over the last two decades, robust optimization (RO) has emerged as a popular approach for decision-making under uncertainty in *single-stage* settings, see e.g., Ben-Tal et al. (2009), Ben-Tal and Nemirovski (2000, 1999, 1998), Bertsimas et al. (2004), Bertsimas and Sim (2004). For example, it has been used successfully to address problems in inventory management (Ardestani-Jaafari and Delage (2016)), network optimization (Bertsimas and Sim (2003)), product pricing (Adida and Perakis (2006), Thiele (2009)), portfolio optimization (Goldfarb and Iyengar (2004, 2003)), and healthcare (Gupta et al. (2017), Bandi et al. (2018), Chan et al. (2018)).

More recently, the robust optimization paradigm has also proved to be an extremely effective means of addressing decision-making problems affected by uncertainty in *two-* and *multi-stage* settings, see e.g., Ben-Tal et al. (2004), Bertsimas et al. (2011), Zhen et al. (2016), Vayanos et al. (2012), Bertsimas and Goyal (2012), Xu and Burer (2018). In these 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* nature of the decision-making process. On the other hand, the requirement that decisions be constant in those parameters that remain unobserved at the time of decision-making captures the *non-anticipative* nature of the decision-process. Two- and multi-stage robust models and solution approaches have proved attractive to address large scale decision-making problems over time. For example, they have been used to successfully tackle sequential problems in energy (Zhao et al. (2013), Jiang et al. (2014)), inventory and supply-chain management (Ben-Tal et al. (2005), Mamani et al. (2017)), network optimization (Atamtürk and Zhang (2007)), vehicle routing (Gounaris et al. (2013)), and process scheduling (Lappas and Gounaris (2016)). For in-depth reviews of the literature on robust optimization and its applications, we refer the reader to Bertsimas et al. (2010), Gabrel et al. (2014), Gorissen et al. (2015), Yanikoglu et al. (2017), Georghiou et al. (2018), and to the references there-in.

Most of the models and solution approaches in two- and multi-stage 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 non-anticipativity requirements depend upon these decisions, severely complicating solution of such problems.

1.1.1. Traditional Application Areas. Over the last three decades, researchers in stochastic programming and robust optimization have investigated several decision-making problems affected by uncertain parameters whose time of revelation is decision-dependent. We detail some of these in the following.

Offshore Oilfield Exploitation. Offshore oilfields consist of several reservoirs of oil whose volume and initial deliverability (maximum initial extraction rate) are uncertain, see e.g., 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.

R&D Project Portfolio Optimization. Research and development firms typically maintain long pipelines of candidate projects whose returns are uncertain, see e.g., 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, the project start times and the 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 e.g., 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.

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 e.g. 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.

Improving Parameter Estimates. In decision-making under uncertainty, the uncertain parameters in the problem often have characteristics (e.g., mean) that are not precisely known, see e.g., Artstein and Wets (1993). If such uncertainties are due to the estimation procedure itself or to the approach used to gather the data, they can be mitigated for example by running a more elaborate mathematical model to improve estimates or by gathering additional or more accurate data. Thus, the decisions to gather additional data and to employ a more sophisticated estimation procedure control information discovery in this problem.

1.1.2. Novel Application Areas. Decision-making problems affected by uncertain parameters whose time of revelation is decision-dependent also arise in a variety of other applications that have received little or no attention in the stochastic and robust optimization literature to date.

Physical Network Monitoring. Physical networks (such as road traffic networks, water pipe networks, or city pavements) are subject to unpredictable disruptions (e.g., road accidents, or physical damage). To help anticipate and resolve such disruptions, static and/or mobile sensors can be strategically positioned to monitor the state of the network. The state of a particular node or arc in the network at a given time is observable only if a sensor is located in its neighborhood at that time. Thus, the sensor positions control the time of information discovery in this problem.

Algorithmic Social Interventions. Algorithmic social interventions rely on social network information to strategically conduct social interventions, e.g., to decide who to train as "peer leaders" in a social network to most effectively spread information about HIV prevention (Wilder et al. (2017)), to decide who to train as "gatekeepers" in a social network to be able to identify warning signs of suicide among their peers, or to select individuals in charge of watching out for their peers during a landslide (Rahmattalabi et al. (2019)). In these applications, the social network of the individuals involved is typically uncertain and significant capital outlays must be made to fully uncover all social ties. Thus, the decisions to query nodes about their social ties control the time of information discovery in this problem.

Active Learning in Machine Learning. In active learning, unlabeled data is usually abundant but manually labeling it is expensive. A learning algorithm can interactively query a user (or other information source such as workers on Amazon Mechanical Turk[1]) to manually label the data. For each available unlabeled data point, the corresponding label will only be revealed if the algorithm chooses to query the user. Thus, the decisions to query a user about a particular unlabeled data point 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 very large set of questions before making a recommendation for a particular item (or a set of items) for purchase. These questions can ask users to quantify how much they like an item or they can take the form of pairwise comparisons between items, see e.g., McElfresh et al. (2019). The answers to the questions are initially unknown and will only be revealed if the particular question is asked. The decisions to ask particular questions thus 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, i.e., to approximate them by here-and-now decisions, and to solve the multi-stage 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 non-anticipativity constraints for the measurement variables.

Our paper most closely relates to the work of Vayanos et al. (2011), where-in the authors investigate two- and multi-stage 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 pre-partitioning of the support of the uncertain parameters. Since this solution approach applies to the class of problems we investigate in this paper, 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 realization of the uncertain parameters and have been investigated by Spacey et al. (2012), Nohadani and Sharma (2016), Nohadani and Roy (2017), Zhang et al. (2017), Bertsimas and Vayanos (2017). The models and solution approaches in these papers do not apply to our setting since, in problems with decision-dependent information discovery, the decision-maker cannot influence the set of possible realization of the uncertain parameters but rather the *information* available about the uncertain parameters. In particular, the problems investigated by Spacey et al. (2012), Nohadani and Sharma (2016), and Nohadani and Roy (2017) are all single-stage (i.e., static) robust problems with decision-dependent uncertainty sets, while problems with decision-dependent information discovery are inherently sequential in nature–indeed, gathering information is only useful if we can use that information to improve our decisions in the future.

Robust Optimization with Binary Adaptive Variables. Two-stage, and to a lesser extent also multi-stage, 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 work investigates the so-called K-adaptability counterpart of two-stage problems, see Bertsimas and Caramanis (2010), Hanasusanto et al. (2015), Subramanyam et al. (2017), 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). Bertsimas and Caramanis (2010) provide necessary conditions for the K-adaptability problem to improve upon the static formulation where all decision are taken here-and-now ($K = 1$) and propose a reformulation of the 2-adaptability problem as a finite-dimensional bilinear problem. Hanasusanto et al. (2015) characterize the complexity of two-stage robust programs for the case where the recourse decisions are binary in terms of the number of second-stage policies K needed to recover the original two-stage robust problem. They also derive explicit mixed-binary linear program (MBLP) reformulations for the K-adaptability problem with objective and constraint uncertainty.

Worst-Case Regret Optimization. Finally, our work relates to two-stage worst-case absolute regret minimization problems. These have received a lot of attention in the last decade as they are often seen as being less conservative than their utility maximizing counterparts, see e.g., Assavapokee et al. (2008b,a), Zhang (2011), Jiang et al. (2013), Ng (2013), Chen et al. (2014), Ning and You (2018), Poursoltani and Delage (2019), and the references therein. To the best of our knowledge, our paper is the first to investigate worst-case regret minimization problems in the presence of endogenous uncertain parameters, and existing approaches cannot be readily applied 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 propose a novel formulation of two- and multi-stage robust problems with decision-dependent information discovery and prove correctness of this formulation. We leverage this model to generalize the K-adaptability approximation approach from the literature to problems with decision-dependent information discovery. This approximation allows us to control the trade-off between complexity and solution quality by tuning a single design parameter, K.

(b) We propose tractable reformulations of the K-adaptability counterpart of problems with decision-dependent information discovery in the form of moderately sized mixed-binary linear programs solvable with off-the shelf solvers. We show that our reformulations subsume as special cases the formulations from the literature that apply only to problems with exogenous uncertain parameters.

(c) We generalize the K-adaptability approximation scheme to problems with piecewise linear convex objective function. We propose a “column-and-constraint” generation algorithm that leverages the decomposable structure of the problem to solve it efficiently. This enables us to address a special class

of problems of practical interest that seek to minimize worst-case absolute regret, i.e., the difference between the worst-case performance of the decision implemented and performance of the best possible decision in hindsight. This generalization and algorithm apply also to problems with exogenous uncertain parameters.

(d) We generalize the K-adaptability approximation approach to multi-stage problems. We also propose a conservative approximation to the K-adaptability counterpart of problems involving continuous recourse decisions. These generalizations apply also to problems with exogenous uncertain parameters.

(e) We propose two novel mathematical formulations of the robust active preference learning problem. We show, both by means of stylized examples and through computational results on randomly generated instances that our proposed approach outperforms the state-of-the-art in the literature in terms of solution time, solution quality, and usability.

(f) We perform a case study based on real data from the Homeless Management Information System[2] (HMIS) to recommend policies that meet the needs of policy-makers at the Los Angeles Homeless Services Authority[3] (LAHSA), the lead agency in charge of allocating public housing resources in L.A. County to those experiencing homelessness. We demonstrate competitive performance relative to the state of the art in terms of solution time, solution quality, and usability. Our case study also highlights the benefits of minimizing worst-case regret relative to maximizing worst-case utility.

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. Sections 4 and 5 propose reformulations of the K-adaptability counterparts of problems with DDID as MBLPs, for problems with objective and constraint uncertainty, respectively. Section 6 generalizes the K-adaptability approximation to problems with piecewise linear convex objective and to the minimization of worst-case regret. Section 7 generalizes the K-adaptability approximation to multi-stage problems. Speed-up strategies and extensions are discussed in Section 8. Section 9 introduces the preference elicitation problem at LAHSA and formulates it as two-stage robust problem with decision-dependent information discovery. Finally, Section 10 discusses our numerical results on both synthetic and real data from the HMIS. The proofs of all statements can be found in the Electronic Companion to the paper.

Notation. Throughout this paper, vectors (matrices) are denoted by boldface lowercase (uppercase) letters. The kth element of a vector $\boldsymbol{x} \in \mathbb{R}^n$ $(k \leq n)$ is denoted by $\boldsymbol{x}_k$. Scalars are denoted by lowercase letters, e.g., α or u. For a matrix $\boldsymbol{H} \in \mathbb{R}^{n \times m}$, we let $[\boldsymbol{H}]_k \in \mathbb{R}^m$ denote the kth 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, i.e., their element-wise product. 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 novel formulation from Section 3, we begin by introducing 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 Ξ is a non-empty bounded polyhedron expressible as $\Xi := \{\boldsymbol{\xi} \in \mathbb{R}^{N_\xi} : \boldsymbol{A}\boldsymbol{\xi} \leq \boldsymbol{b}\}$ 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 for $\boldsymbol{y}$ to adjust to the realization of $\boldsymbol{\xi}$ is modeled.

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 $\boldsymbol{\xi}$. Mathematically, $\boldsymbol{y}$ is modeled as a function (or *decision rule*) of $\boldsymbol{\xi}$ that is selected today, along with $\boldsymbol{x}$. Under this modeling paradigm, a generic two-stage linear robust problem with exogenous uncertainty is expressible as:

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

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 $\boldsymbol{\xi}_{N_\xi+1}$ and augmenting the uncertainty set with the constraint $\boldsymbol{\xi}_{N_\xi+1}=1$.

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 modeling paradigm, a generic two-stage robust problem with exogenous uncertainty is expressible as:

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

Problems (1) and (2) are equivalent in a sense made formal in the following theorem.

Theorem 1. *The optimal objective values of Problems* (1) *and* (2) *are equal. Moreover, the following statements hold true:*

(i) Let $\boldsymbol{x}$ *be optimal in Problem* (2) *and, for each* $\boldsymbol{\xi}\in\Xi$*, let*

$$
\boldsymbol{y}(\boldsymbol{\xi}) \in \operatorname*{arg\,min}_{\boldsymbol{y}\in\mathcal{Y}} \; \left\{\boldsymbol{\xi}^\top \boldsymbol{C}\,\boldsymbol{x} + \boldsymbol{\xi}^\top \boldsymbol{Q}\,\boldsymbol{y} \,:\, \boldsymbol{T}\boldsymbol{x}+\boldsymbol{W}\boldsymbol{y}\le \boldsymbol{H}\boldsymbol{\xi}\right\}.
$$

Then, $(\boldsymbol{x},\boldsymbol{y}(\cdot))$ *is optimal in Problem* (1)*.*

(ii) Let $(\boldsymbol{x},\boldsymbol{y}(\cdot))$ *be optimal in Problem* (1)*. Then,* $\boldsymbol{x}$ *is optimal in Problem* (2)*.*

The theorem above is, to some extent, well known in the literature, see e.g., Shapiro (2017). We provide a proof in the Electronic Companion EC.2 to keep the paper self contained.

Remark 1. *We note that the results in Sections 2 and 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.*

While the models above are equivalent, each of them has proved successful in different contexts. Specifically, Problem (1) has been the underpinning building block of most of the literature on the decision rule approximation for problems with both continuous and discrete wait-and-see decisions, see Section 1 for references. Problem (2) on the other hand 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 e.g., Zeng and Zhao (2013), Ayoub and Poss (2016), and Bertsimas and Shtern (2017).

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

In this section, we propose a novel formulation of two-stage robust optimization problems with *decision-dependent information discovery*. This formulation underpins our ability to generalize the popular K-adaptability approximation approach from the literature to two-stage problems affected by uncertain parameters whose time of revelation is decision-dependent, see Sections 4 and 5. This section is organized as follows. In Section 3.1, we describe two-stage robust optimization problems with DDID. We present a formulation of a generic two-stage robust optimization problem with decision-dependent information discovery from the literature and an associated approximate solution scheme in Sections 3.2 and 3.3, respectively. Then, in Section 3.4, we propose an equivalent, novel formulation, that constitutes the main result of this section. Finally, in Section 3.5, we introduce the K-adaptability approximation scheme for problems with DDID.

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 ith 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, $\boldsymbol{y}$ should 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 *non-anticipativity* in the literature. In the presence of uncertain parameters whose time of revelation is decision-dependent, this requirement translates to *decision-dependent non-anticipativity constraints*. In addition, the decisions $\boldsymbol{w}$ now impact both the objective function and the 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 $\boldsymbol{w}_i = 0$, 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

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

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 $\boldsymbol{w}_i \le \boldsymbol{w}_{i'}$. Accordingly, it can postulate that the total number of uncertain parameters that are observed does not exceed a certain budget Q using $\sum_{i=1}^{N_\xi} \boldsymbol{w}_i \le Q$. The set $\mathcal{W}$ can also be used to capture exogenous uncertain parameters. Indeed, if uncertain parameter $\boldsymbol{\xi}_i$ is exogenous (i.e., automatically observed between the first and second decision-stages), we require $\boldsymbol{w}_i = 1$. On the other hand, if an uncertain parameter can only be observed after the decision $\boldsymbol{y}$ is made, then we require $\boldsymbol{w}_i = 0$. The last constraint in the problem is a decision-dependent non-anticipativity 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 indistinguishable. 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.

3.3. Decision Rule based Approximation Approach from the Literature

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 pre-selected partition of the uncertainty set. They partition Ξ into hyper-rectangles of the form

$$\Xi_{\boldsymbol{s}} := \left\{ \boldsymbol{\xi} \in \Xi : \boldsymbol{c}^i_{\boldsymbol{s}_i - 1} \leq \boldsymbol{\xi}_i < \boldsymbol{c}^i_{\boldsymbol{s}_i},\ i = 1, \ldots, k \right\},$$

where $\boldsymbol{s} \in \mathcal{S} := \times_{i=1}^{N_\xi} \{1, \ldots, \boldsymbol{r}_i\} \subseteq \mathbb{Z}^{N_\xi}$ and

$$\boldsymbol{c}^i_1 < \boldsymbol{c}^i_2 < \cdots < \boldsymbol{c}^i_{\boldsymbol{r}_i - 1} \quad \text{for } i = 1, \ldots, N_\xi$$

represent $\boldsymbol{r}_i - 1$ breakpoints along the $\boldsymbol{\xi}_i$ axis. They approximate binary decision rules $\boldsymbol{y}_i$ such that $\boldsymbol{y}_i(\boldsymbol{\xi}) \in \{0,1\}\ \forall \boldsymbol{\xi} \in \Xi$, in Problem (3), by piecewise constant decision rules of the form

$$\boldsymbol{y}_i(\boldsymbol{\xi}) = \sum_{\boldsymbol{s} \in \mathcal{S}} \mathbb{I}\left(\boldsymbol{\xi} \in \Xi_{\boldsymbol{s}}\right) \boldsymbol{y}^{\boldsymbol{s}}_i$$

for some $\boldsymbol{y}^{\boldsymbol{s}}_i \in \{0,1\}$, $\boldsymbol{s} \in \mathcal{S}$. Similarly, they approximate the real-valued decisions $\boldsymbol{y}_i$ such that $\boldsymbol{y}_i(\boldsymbol{\xi}) \in \mathbb{R}$, in Problem (3), by piecewise linear decision rules of the form

$$\boldsymbol{y}_i(\boldsymbol{\xi}) = \sum_{\boldsymbol{s} \in \mathcal{S}} \mathbb{I}\left(\boldsymbol{\xi} \in \Xi_{\boldsymbol{s}}\right) (\boldsymbol{y}^{\boldsymbol{s},i})^\top \boldsymbol{\xi}$$

for some $\boldsymbol{y}^{\boldsymbol{s},i} \in \mathbb{R}^{N_\xi}$, $\boldsymbol{s} \in \mathcal{S}$. They show that, under this representation, the decision-dependent non-anticipativity constraints in Problem (3) are expressible as

$$\begin{array}{lll}
|\boldsymbol{y}^{\boldsymbol{s}}_i - \boldsymbol{y}^{\boldsymbol{s}'}_i| \leq \boldsymbol{w}_j & \forall j \in \{1, \ldots, N_\xi\},\ i \in \{1, \ldots, N_y\} : \boldsymbol{y}_i(\boldsymbol{\xi}) \in \{0,1\}\ \forall \boldsymbol{\xi},\ \forall \boldsymbol{s},\ \boldsymbol{s}' \in \mathcal{S} : \boldsymbol{s}_{-j} = \boldsymbol{s}'_{-j} \\
|\boldsymbol{y}^{\boldsymbol{s},i}_{j'} - \boldsymbol{y}^{\boldsymbol{s}',i}_{j'}| \leq M \boldsymbol{w}_j & \forall j, j' \in \{1, \ldots, N_\xi\},\ i \in \{1, \ldots, N_y\} : \boldsymbol{y}_i(\boldsymbol{\xi}) \in \mathbb{R}\ \forall \boldsymbol{\xi},\ \forall \boldsymbol{s},\ \boldsymbol{s}' \in \mathcal{S} : \boldsymbol{s}_{-j} = \boldsymbol{s}'_{-j} \\
|\boldsymbol{y}^{\boldsymbol{s},i}_j| \leq M \boldsymbol{w}_j & \forall j \in \{1, \ldots, N_\xi\},\ i \in \{1, \ldots, N_y\} : \boldsymbol{y}_i(\boldsymbol{\xi}) \in \mathbb{R}\ \forall \boldsymbol{\xi},\ \forall \boldsymbol{s} \in \mathcal{S}.
\end{array}$$

The first two sets of constraints impose non-anticipativity across distinct subsets subsets of the partition for the binary and continuous valued decisions, respectively. The last set of constraints imposes non-anticipativity for the linear decision rules within each subset. We henceforth refer to this decision rule approximation approach from the literature as the "prepartitioning" approach.

Unfortunately, as the following example illustrates, this approach is highly sensitive to the choice in the breakpoint configuration.

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

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

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* $\boldsymbol{\xi}_i > \boldsymbol{\epsilon}_i$ *(resp.* $\boldsymbol{\xi}_i < \boldsymbol{\epsilon}_i$*). Thus, from the decision-dependent non-anticipativity constraints, the only feasible choice for* $\boldsymbol{w}$ *is* $\mathbf{e}$.

Consider partitioning Ξ *according to the approach from Vayanos et al. (2011). If, for all* $i \in \{1,2\}$*, there exists* $j \in \{1,\ldots,\boldsymbol{r}_i - 1\}$ *such that* $\boldsymbol{c}^i_j = \boldsymbol{\epsilon}_i$*, then the pre-partitioning approach yields an optimal solution to Problem* (4) *with objective value 0. On the other hand, if for some* $i \in \{1,2\}$*, the* j*th breakpoint satisfies* $\boldsymbol{c}^i_j \neq \boldsymbol{\epsilon}_i$ *for all* $j \in \{1,\ldots,\boldsymbol{r}_i - 1\}$*, then the approach from Vayanos et al. (2011) yields an infeasible problem and the optimality gap of the pre-partitioning approach is infinite. In particular, suppose* $\boldsymbol{\epsilon} = 1\text{e}{-3}\mathbf{e}$ *and that we uniformly partition each axis iteratively in 2, 3, 4, etc. subsets. Then, we will need to introduce 1999 breakpoints along each direction before reaching a feasible (and thus optimal) solution, giving a total of* 4e7 *subsets and a problem formulation involving over* 8e7 *binary decision variables and over* 16e7 *constraints.*

Example 1 is not surprising: the approach from Vayanos et al. (2011) was motivated by stochastic programming problems which are less sensitive to the breakpoint configuration than robust problems. As we will see in Section 10, our proposed approach outperforms that presented in Vayanos et al. (2011) in the case of *robust* problems with DDID, in terms of both solution time and solution quality.

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

3.4. 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 propose to derive an analogous *dynamic* formulation for the case of endogenous uncertainties. In particular, we propose to 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 Ξ. 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 she has not observed, i.e., for which $\boldsymbol{w}_i = 0$. Indeed, the decision $\boldsymbol{y}$ may have to be taken under uncertainty if there is some i such that $\boldsymbol{w}_i = 0$, in which case not all of the uncertain parameters have been revealed when $\boldsymbol{y}$ is selected. Indeed, after $\boldsymbol{y}$ is selected, nature is free to choose any realization of $\boldsymbol{\xi} \in \Xi$ that is compatible with the original choice $\overline{\boldsymbol{\xi}}$ in the sense that $\boldsymbol{\xi}_i = \overline{\boldsymbol{\xi}}_i$ for all i such that $\boldsymbol{w}_i = 1$. This model captures the notion that, after $\boldsymbol{y}$ has been selected, nature is still free to choose the elements $\boldsymbol{\xi}_i$ that have not been observed (i.e., for which $\boldsymbol{w}_i = 0$) provided it does so in a way that is consistent with those parameters that *have* been observed. Mathematically, given the measurement decisions $\boldsymbol{w}$ and the observation $\overline{\boldsymbol{\xi}}$, nature can select any element $\boldsymbol{\xi}$ from the set

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

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

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

Note that, at the time when $\boldsymbol{y}$ is selected, some elements of $\boldsymbol{\xi}$ are still uncertain. The choice of $\boldsymbol{y}$ thus needs to be robust to the choice of those uncertain parameters that remain to be revealed. In particular, the

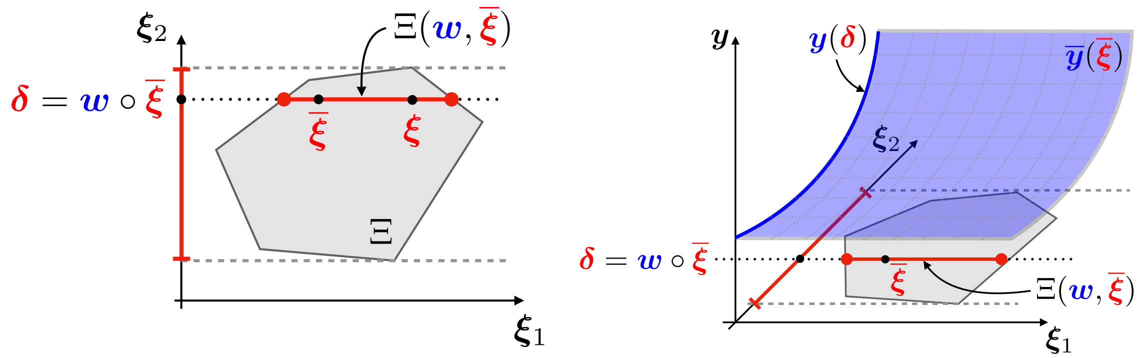

Figure 1 The figure on the left illustrates the role played by $\overline{\boldsymbol{\xi}}$ in the new formulation $(\mathcal{P})$ and the definition of the uncertainty sets Ξ and $\Xi(\boldsymbol{w},\overline{\boldsymbol{\xi}})$. Consider a setting where $\Xi \subseteq \mathbb{R}^2$ (i.e., $N_\xi = 2$) and suppose that $\boldsymbol{w} = (0,1)$ so that the decision-maker has chosen to only observe $\boldsymbol{\xi}_2$. In the figures, Ξ is shown as the grey shaded area. Once $\overline{\boldsymbol{\xi}}$ is chosen by nature, the decision-maker can only infer that $\boldsymbol{\xi}$ will materialize in the set $\Xi(\boldsymbol{w},\overline{\boldsymbol{\xi}})$ which collects all parameter realizations $\boldsymbol{\xi} \in \Xi$ that satisfy $\boldsymbol{\xi}_2 = \overline{\boldsymbol{\xi}}_2$, being compatible with our partial observation. The figure on the right illustrates the construction of an optimal non-anticipative decision $\overline{\boldsymbol{y}}$ from the an optimal solution $\boldsymbol{y}(\boldsymbol{\delta})$ to $\min_{\boldsymbol{y} \in \mathcal{Y}} \left\{ \max_{\boldsymbol{\xi} \in \Xi(\boldsymbol{w},\boldsymbol{\delta})} \; \boldsymbol{\xi}^\top \boldsymbol{C}\,\boldsymbol{x} + \boldsymbol{\xi}^\top \boldsymbol{D}\,\boldsymbol{w} + \boldsymbol{\xi}^\top \boldsymbol{Q}\,\boldsymbol{y} \,:\, \boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y} \le \boldsymbol{H}\boldsymbol{\xi} \quad \forall \boldsymbol{\xi} \in \Xi(\boldsymbol{w},\boldsymbol{\delta}) \right\}$, see Theorem 2. We note that the policy $\overline{\boldsymbol{y}}$ constructed as in Theorem 2 is constant along the $\boldsymbol{\xi}_1$ direction since here $\boldsymbol{w}_1 = 0$.

constraints need to be satisfied for all choices of $\boldsymbol{\xi} \in \Xi(\boldsymbol{w},\overline{\boldsymbol{\xi}})$. Accordingly, $\boldsymbol{y}$ is chosen so as to minimize the worst-case possible cost when $\boldsymbol{\xi}$ is valued in the set $\boldsymbol{\xi} \in \Xi(\boldsymbol{w},\overline{\boldsymbol{\xi}})$.

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

Theorem 2. *The optimal objective values of Problems* (3) *and* $(\mathcal{P})$ *are equal. Moreover, the following statements hold true:*

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

$$\boldsymbol{y}'(\boldsymbol{\delta}) \in \operatorname*{arg\,min}_{\boldsymbol{y} \in \mathcal{Y}} \left\{ \max_{\boldsymbol{\xi} \in \Xi(\boldsymbol{w},\boldsymbol{\delta})} \; \boldsymbol{\xi}^\top \boldsymbol{C}\,\boldsymbol{x} + \boldsymbol{\xi}^\top \boldsymbol{D}\,\boldsymbol{w} + \boldsymbol{\xi}^\top \boldsymbol{Q}\,\boldsymbol{y} \,:\, \boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y} \le \boldsymbol{H}\boldsymbol{\xi} \quad \forall \boldsymbol{\xi} \in \Xi(\boldsymbol{w},\boldsymbol{\delta}) \right\}.$$

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

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

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

$$\min_{\boldsymbol{y}\in\mathcal{Y}}\left\{\max_{\boldsymbol{\xi}\in\Xi(\boldsymbol{w},\overline{\boldsymbol{\xi}})}\ \boldsymbol{\xi}^\top\boldsymbol{C}\,\boldsymbol{x}+\boldsymbol{\xi}^\top\boldsymbol{D}\,\boldsymbol{w}+\boldsymbol{\xi}^\top\boldsymbol{Q}\,\boldsymbol{y}\ :\ \boldsymbol{T}\boldsymbol{x}+\boldsymbol{V}\boldsymbol{w}+\boldsymbol{W}\boldsymbol{y}\le\boldsymbol{H}\boldsymbol{\xi}\quad\forall\boldsymbol{\xi}\in\Xi(\boldsymbol{w},\overline{\boldsymbol{\xi}})\right\},$$

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

$$\tilde{\boldsymbol{y}}(\overline{\boldsymbol{\xi}})\in\operatorname*{arg\,min}_{\boldsymbol{y}\in\mathcal{Y}}\left\{\max_{\boldsymbol{\xi}\in\Xi(\boldsymbol{w},\overline{\boldsymbol{\xi}})}\ \boldsymbol{\xi}^\top\boldsymbol{C}\,\boldsymbol{x}+\boldsymbol{\xi}^\top\boldsymbol{D}\,\boldsymbol{w}+\boldsymbol{\xi}^\top\boldsymbol{Q}\,\boldsymbol{y}\ :\ \boldsymbol{T}\boldsymbol{x}+\boldsymbol{V}\boldsymbol{w}+\boldsymbol{W}\boldsymbol{y}\le\boldsymbol{H}\boldsymbol{\xi}\quad\forall\boldsymbol{\xi}\in\Xi(\boldsymbol{w},\overline{\boldsymbol{\xi}})\right\},$$

may not be constant in those parameters that remain unobserved. We note of course that other tie-breaking mechanisms could be used to build a non-anticipative 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. Before introducing the K-adaptability approximation, we investigate a concrete instance of Problem ($\mathcal{P}$) consisting of binary recourse decisions only.

Example 2. *Consider the following instance of Problem* ($\mathcal{P}$)*, adapted from Hanasusanto et al. (2015) to incorporate decision-dependent information discovery.*

$$\begin{array}{llll}\underset{\boldsymbol{w}\in\{0,1\}^2}{\text{minimize}} & \max\limits_{\overline{\boldsymbol{\xi}}\in\Xi} & \min\limits_{\boldsymbol{y}\in\{0,1\}^2} & \max\limits_{\boldsymbol{\xi}\in\Xi(\boldsymbol{w},\overline{\boldsymbol{\xi}})}\ (\boldsymbol{\xi}_1+\boldsymbol{\xi}_2)(\boldsymbol{y}_2-\boldsymbol{y}_1)+\boldsymbol{d}_1\boldsymbol{w}_1+\boldsymbol{d}_2\boldsymbol{w}_2\\ & & \text{s.t.} & \boldsymbol{y}_1\ge\boldsymbol{\xi}_1\quad\forall\boldsymbol{\xi}\in\Xi(\boldsymbol{w},\overline{\boldsymbol{\xi}})\\ & & & \boldsymbol{y}_1+\boldsymbol{y}_2=1,\end{array}\tag{5}$$

where $\boldsymbol{d}_1$, $\boldsymbol{d}_2\in(0,1)$ *are given scalars representing the observation costs associated with* $\boldsymbol{w}_1$ *and* $\boldsymbol{w}_2$*, respectively, and* $\Xi:=\{\boldsymbol{\xi}\in\mathbb{R}^2\ :\ -1\le\boldsymbol{\xi}_1\le1,\ -1.1\le\boldsymbol{\xi}_2\le1\}$*. For each feasible choice of* $\boldsymbol{w}$*, we investigate the associated optimal wait-and-see decision, as well as the corresponding objective function value, see Figure 2.*

Consider the choice $\boldsymbol{w}=\boldsymbol{0}$*, whereby no uncertain parameter is observed between the first and second decision stages. Then,* $\Xi(\boldsymbol{w},\overline{\boldsymbol{\xi}})=\Xi$*. Under this here-and-now decision, Problem* (5) *is expressible as a single-stage robust problem as follows*

$$\left\{\min_{\boldsymbol{y}\in\{0,1\}^2}\ \left(\max_{\boldsymbol{\xi}\in\Xi}\ (\boldsymbol{\xi}_1+\boldsymbol{\xi}_2)(\boldsymbol{y}_2-\boldsymbol{y}_1)\right)\ :\ \boldsymbol{y}_1\ge\boldsymbol{\xi}_1\ \forall\boldsymbol{\xi}\in\Xi,\ \boldsymbol{y}_1+\boldsymbol{y}_2=1\right\}.$$

It can can be readily verified that the only feasible (and therefore optimal) wait-and-see action in this case is $\boldsymbol{y}=(1,0)$*, a static decision. The associated objective function and corresponding value is*

$$\max_{\boldsymbol{\xi}\in\Xi}\ -(\boldsymbol{\xi}_1+\boldsymbol{\xi}_2)\ =\ 2.1.$$

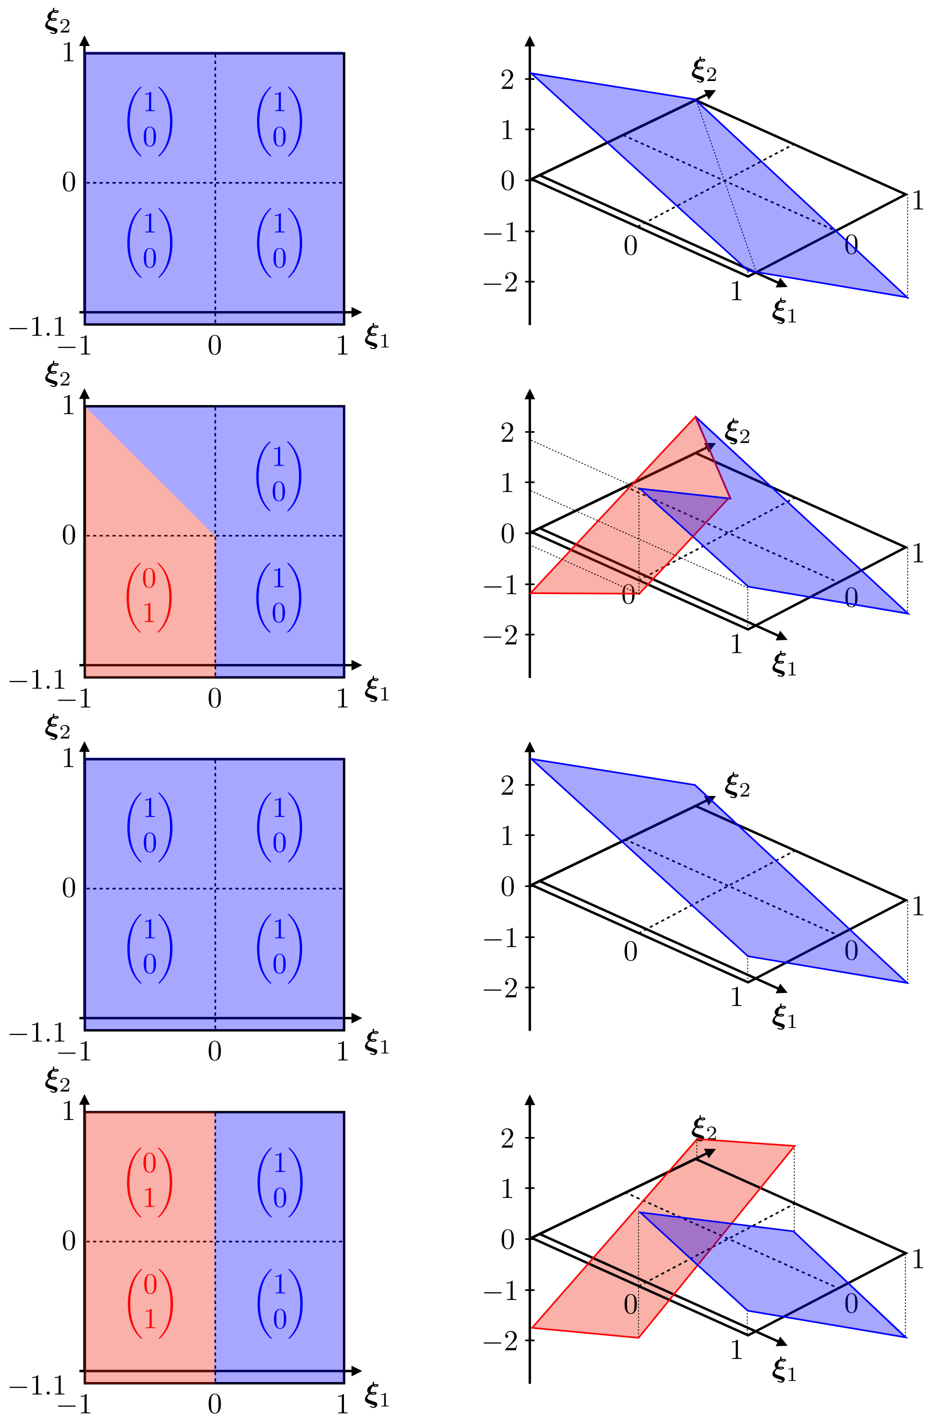

Figure 2 Companion figure for Example 2, assuming $d_1 = d_2 = 0.4$. Optimal wait-and-see decision $\boldsymbol{y}$ (left) and associated objective function (right) for the cases when $\boldsymbol{w} = \mathbf{0}$ (first row), $\boldsymbol{w} = \mathbf{e}$ (second row), $\boldsymbol{w} = (0,1)$ (third row), and $\boldsymbol{w} = (1,0)$ (last row) in Problem (5). The optimal solution is given by $\boldsymbol{w}^\star = (1,0)$. For the optimal solution $\boldsymbol{w}^\star$, the objective function is discontinuous on the set $\{\boldsymbol{\xi} \in \Xi : \boldsymbol{\xi}_1 = 0\}$ and in particular the optimal objective value is not attained.

Consider the choice $\boldsymbol{w} = \mathbf{e}$, whereby both uncertain parameters are observed between the first and second decision stages. Then, $\Xi(\boldsymbol{w}, \overline{\boldsymbol{\xi}}) = \{\overline{\boldsymbol{\xi}}\}$. Under this here-and-now decision, Problem (5) *reduces to*

$$\max_{\boldsymbol{\xi} \in \Xi} \left\{ \min_{\boldsymbol{y} \in \{0,1\}^2} \;\; (\boldsymbol{\xi}_1 + \boldsymbol{\xi}_2)(\boldsymbol{y}_2 - \boldsymbol{y}_1) + \boldsymbol{d}_1 + \boldsymbol{d}_2 \;\; : \;\; \boldsymbol{y}_1 \geq \boldsymbol{\xi}_1, \;\; \boldsymbol{y}_1 + \boldsymbol{y}_2 = 1 \right\}.$$

The constraints in the problem imply that $\boldsymbol{y} = (1,0)$ is the only feasible (and therefore optimal) solution whenever $\boldsymbol{\xi}_1 > 0$. For $\boldsymbol{\xi}_1 \leq 0$, the optimal choices are $\boldsymbol{y} = (1,0)$ if $\boldsymbol{\xi}_1 + \boldsymbol{\xi}_2 \geq 0$, and $\boldsymbol{y} = (0,1)$, else. The associated objective function is

$$\boldsymbol{d}_1 + \boldsymbol{d}_2 + \max_{\boldsymbol{\xi} \in \Xi} \begin{cases} -(\boldsymbol{\xi}_1 + \boldsymbol{\xi}_2) & \textit{if } (\boldsymbol{\xi}_1 > 0) \textit{ or } (\boldsymbol{\xi}_1 \leq 0 \textit{ and } \boldsymbol{\xi}_1 + \boldsymbol{\xi}_2 \geq 0) \\ \boldsymbol{\xi}_1 + \boldsymbol{\xi}_2 & \textit{else,} \end{cases}$$

yielding an objective value of $(1.1 + \boldsymbol{d}_1 + \boldsymbol{d}_2)$ that is not attained.

Consider the choice $\boldsymbol{w} = (0,1)$, whereby only $\boldsymbol{\xi}_2$ is observed between the first and second decision stages. Then, Problem (5) *reduces to*

$$\max_{\boldsymbol{\xi}_2 \in [-1.1,1]} \left\{ \min_{\boldsymbol{y} \in \{0,1\}^2} \;\; \max_{\boldsymbol{\xi}_1 \in [-1,1]} \; (\boldsymbol{\xi}_1 + \boldsymbol{\xi}_2)(\boldsymbol{y}_2 - \boldsymbol{y}_1) + \boldsymbol{d}_2 \;\; : \;\; \boldsymbol{y}_1 \geq \boldsymbol{\xi}_1 \quad \forall \boldsymbol{\xi}_1 \in [-1,1], \;\; \boldsymbol{y}_1 + \boldsymbol{y}_2 = 1 \right\}.$$

For any choice of $\boldsymbol{\xi}_2$, the only option for the wait-and-see decision is $\boldsymbol{y} = (1,0)$ (since $\boldsymbol{\xi}_1$ remains uncertain). The associated objective function and corresponding objective value is

$$\boldsymbol{d}_2 + \max \;\; \{-(\boldsymbol{\xi}_1 + \boldsymbol{\xi}_2) \, : \, \boldsymbol{\xi}_1 \in [-1,1], \, \boldsymbol{\xi}_2 \in [-1.1,1]\} \; = \; \boldsymbol{d}_2 + 2.1.$$

Lastly, consider the choice $\boldsymbol{w} = (1,0)$, whereby only $\boldsymbol{\xi}_1$ is observed between the first and second decision stages. Then, Problem (5) *reduces to*

$$\max_{\boldsymbol{\xi}_1 \in [-1,1]} \left\{ \min_{\boldsymbol{y} \in \{0,1\}^2} \;\; \max_{\boldsymbol{\xi}_2 \in [-1.1,1]} \; (\boldsymbol{\xi}_1 + \boldsymbol{\xi}_2)(\boldsymbol{y}_2 - \boldsymbol{y}_1) + \boldsymbol{d}_1 \;\; : \;\; \boldsymbol{y}_1 \geq \boldsymbol{\xi}_1, \;\; \boldsymbol{y}_1 + \boldsymbol{y}_2 = 1 \right\}.$$

For $\boldsymbol{\xi}_1 > 0$, the only feasible (and therefore optimal) choice is $\boldsymbol{y} = (1,0)$. If $\boldsymbol{\xi}_1 \leq 0$, then the optimal wait-and-see decision is $\boldsymbol{y} = (0,1)$. The associated objective function is

$$\boldsymbol{d}_1 + \max_{\boldsymbol{\xi} \in \Xi} \begin{cases} -(\boldsymbol{\xi}_1 + \boldsymbol{\xi}_2) & \textit{if } \boldsymbol{\xi}_1 > 0 \\ \boldsymbol{\xi}_1 + \boldsymbol{\xi}_2 & \textit{else,} \end{cases}$$

yielding an objective value of $(1.1 + \boldsymbol{d}_1)$ that is not attained by any feasible solution.

We conclude that, since $\boldsymbol{d}_1$, $\boldsymbol{d}_2 \in (0,1)$, the optimal solution to Problem (5) *is $\boldsymbol{w}^\star = (1,0)$ with associated optimal objective value $(1.1 + \boldsymbol{d}_1)$ which is never attained.*

The above example shows 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 non closed and non-convex 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). Example 2 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 and 9. Yet, as illustrated by the above example, they pose several theoretical and practical challenges. As we will see in the following sections, whether we are or not able to reformulate the problem exactly as an MILP 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 5.

3.5. K-Adaptability for Problems with Decision-Dependent Information Discovery

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

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

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 ($\mathcal{P}$).

The Price of Usability. We note that 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, and 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.

Remark 2. *If $\boldsymbol{w} = \mathbf{e}$, $\boldsymbol{D} = \mathbf{0}$, and $\boldsymbol{V} = \mathbf{0}$, then $\Xi(\boldsymbol{w}, \overline{\boldsymbol{\xi}}) = \{\overline{\boldsymbol{\xi}}\}$ and therefore Problem $(\mathcal{P}_K)$ reduces to*

$$\begin{array}{ll} \text{minimize} & \max_{\boldsymbol{\xi} \in \Xi} \left[\boldsymbol{\xi}^\top \boldsymbol{C}\, \boldsymbol{x} + \min_{k \in \mathcal{K}} \left\{ \boldsymbol{\xi}^\top \boldsymbol{Q}\, \boldsymbol{y}^k \,:\, \boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{H}\boldsymbol{\xi} \right\} \right] \\ \text{subject to} & \boldsymbol{x} \in \mathcal{X},\, \boldsymbol{y}^k \in \mathcal{Y},\, k \in \mathcal{K}, \end{array} \tag{6}$$

the K-adaptability counterpart of Problem (2) *with only exogenous objective uncertainty, originally studied by Bertsimas and Caramanis (2010) and then Hanasusanto et al. (2015).*

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–indeed, we are in the face of a min-max-min-max, rather than simply min-max-min, 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*

$$\begin{array}{ll} \min & \max_{\{\boldsymbol{\xi}^k\}_{k \in \mathcal{K}} \in \Xi^K(\boldsymbol{w})} \; \min_{k \in \mathcal{K}} \left\{ (\boldsymbol{\xi}^k)^\top \boldsymbol{C}\, \boldsymbol{x} + (\boldsymbol{\xi}^k)^\top \boldsymbol{D}\, \boldsymbol{w} + (\boldsymbol{\xi}^k)^\top \boldsymbol{Q}\, \boldsymbol{y}^k \,:\, \boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{H}\boldsymbol{\xi}^k \right\} \\ \text{s.t.} & \boldsymbol{x} \in \mathcal{X},\, \boldsymbol{w} \in \mathcal{W},\, \boldsymbol{y}^k \in \mathcal{Y},\, k \in \mathcal{K}, \end{array} \tag{7}$$

where

$$\Xi^K(\boldsymbol{w}) := \left\{ \{\boldsymbol{\xi}^k\}_{k \in \mathcal{K}} \in \Xi^K \,:\, \exists \overline{\boldsymbol{\xi}} \in \Xi \text{ such that } \boldsymbol{\xi}^k \in \Xi(\boldsymbol{w}, \overline{\boldsymbol{\xi}}) \text{ for all } k \in \mathcal{K} \right\}. \tag{8}$$

For any fixed $\boldsymbol{w} \in \mathcal{W}$, the subvector $\boldsymbol{\xi}^k$ in the definition of $\boldsymbol{\Xi}^K(\boldsymbol{w})$ 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 $\boldsymbol{\Xi}^K(\boldsymbol{w})$ 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 (7), 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. We derive explicit MBLP reformulations for the K-adaptability counterpart ($\mathcal{P}_K$) with objective and constraint uncertainty, see Sections 4 and 5, respectively. We generalize the K-adaptability approximation to problems with piecewise linear convex objective and to multi-stage problems with DDID in Sections 6 and 7, respectively.

4. The K-Adaptability Problem with Objective Uncertainty

4.1. The K-Adaptability Problem

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

$$\begin{array}{ll} \text{minimize} & \max\limits_{\overline{\boldsymbol{\xi}} \in \Xi} \; \min\limits_{\boldsymbol{y} \in \mathcal{Y}} \left\{ \max\limits_{\boldsymbol{\xi} \in \Xi(\boldsymbol{w}, \overline{\boldsymbol{\xi}})} \; \boldsymbol{\xi}^\top \boldsymbol{C}\, \boldsymbol{x} + \boldsymbol{\xi}^\top \boldsymbol{D}\, \boldsymbol{w} + \boldsymbol{\xi}^\top \boldsymbol{Q}\, \boldsymbol{y} \; : \; \boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y} \le \boldsymbol{h} \right\} \\ \text{subject to} & \boldsymbol{x} \in \mathcal{X},\; \boldsymbol{w} \in \mathcal{W}, \end{array} \tag{$\mathcal{PO}$}$$

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

$$\begin{array}{ll} \text{minimize} & \max\limits_{\overline{\boldsymbol{\xi}} \in \Xi} \; \min\limits_{k \in \mathcal{K}} \left\{ \max\limits_{\boldsymbol{\xi} \in \Xi(\boldsymbol{w}, \overline{\boldsymbol{\xi}})} \; \boldsymbol{\xi}^\top \boldsymbol{C}\, \boldsymbol{x} + \boldsymbol{\xi}^\top \boldsymbol{D}\, \boldsymbol{w} + \boldsymbol{\xi}^\top \boldsymbol{Q}\, \boldsymbol{y}^k \; : \; \boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \le \boldsymbol{h} \right\} \\ \text{subject to} & \boldsymbol{x} \in \mathcal{X},\; \boldsymbol{w} \in \mathcal{W},\; \boldsymbol{y}^k \in \mathcal{Y},\; k \in \mathcal{K}. \end{array} \tag{$\mathcal{PO}_K$}$$

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

$$\begin{array}{ll} \text{minimize} & \max\limits_{\{\boldsymbol{\xi}^k\}_{k \in \mathcal{K}} \in \Xi^K(\boldsymbol{w})} \; \min\limits_{k \in \mathcal{K}} \; \left\{ (\boldsymbol{\xi}^k)^\top \boldsymbol{C}\, \boldsymbol{x} + (\boldsymbol{\xi}^k)^\top \boldsymbol{D}\, \boldsymbol{w} + (\boldsymbol{\xi}^k)^\top \boldsymbol{Q}\, \boldsymbol{y}^k \; : \; \boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \le \boldsymbol{h} \right\} \\ \text{subject to} & \boldsymbol{x} \in \mathcal{X},\; \boldsymbol{w} \in \mathcal{W},\; \boldsymbol{y}^k \in \mathcal{Y},\; k \in \mathcal{K}, \end{array} \tag{9}$$

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 lemma.

Lemma 2. *The K-adaptability counterpart of the two-stage robust optimization problem with decision-dependent information discovery, Problem ($\mathcal{PO}_K$), is equivalent to*

$$\begin{array}{ll} \text{minimize} & \max\limits_{\{\boldsymbol{\xi}^k\}_{k \in \mathcal{K}} \in \Xi^K(\boldsymbol{w})} \; \min\limits_{k \in \mathcal{K}} \; \left\{ (\boldsymbol{\xi}^k)^\top \boldsymbol{C}\, \boldsymbol{x} + (\boldsymbol{\xi}^k)^\top \boldsymbol{D}\, \boldsymbol{w} + (\boldsymbol{\xi}^k)^\top \boldsymbol{Q}\, \boldsymbol{y}^k \right\} \\ \text{subject to} & \boldsymbol{x} \in \mathcal{X},\; \boldsymbol{w} \in \mathcal{W},\; \boldsymbol{y}^k \in \mathcal{Y},\; k \in \mathcal{K} \\ & \boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \le \boldsymbol{h} \quad \forall k \in \mathcal{K}, \end{array} \tag{10}$$

where $\Xi^K(\boldsymbol{w})$ is as defined in Equation (8).

Note that for all $\boldsymbol{w} \in \mathcal{W}$, the set $\Xi^K(\boldsymbol{w})$ is non-empty 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 (10) if $\boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{h}$ for all $k \in \mathcal{K}$, whereas to be feasible in Problem (9) (and accordingly in Problem ($\mathcal{PO}_K$)), it need only satisfy $\boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{h}$ for some $k \in \mathcal{K}$. Thus, a triplet $(\boldsymbol{x}, \boldsymbol{w}, \boldsymbol{y}^k)$ feasible in (9) (and thus in ($\mathcal{PO}_K$)) need not be feasible in Problem (10). However, the proof of Lemma 2, provides a concrete way to construct a feasible solution for Problem (10) from a feasible solution to Problem (9) that achieves the same optimal value.

The lemma above will be the key to reformulating Problem ($\mathcal{PO}_K$) as an MBLP, see Section 4.2. It also enables us to analyze the complexity of evaluating the objective function of the K-adaptability problem under a fixed decision. Indeed, from Problem (10), 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 an LP obtained by writing (10) in epigraph form. We formalize this result in the following.

Observation 1. *For any fixed K and decision $(\boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}^k\}_{k\in\mathcal{K}})$, the optimal objective value of the K-adaptability problem ($\mathcal{PO}_K$) can be evaluated in polynomial time in the size of the input.*

4.2. Reformulation as a Mixed-Binary Linear Program

In Observation 1, 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 and linearizing the resulting bilinear terms, we can obtain an equivalent reformulation of Problem ($\mathcal{PO}_K$) in the form of an MBLP.

Theorem 3. *Suppose $\mathcal{X} \subseteq \{0,1\}^{N_x}$. Then, Problem ($\mathcal{PO}_K$) is equivalent to the following MBLP.*

$$
\begin{array}{ll}
\text{minimize} & \boldsymbol{b}^\top \boldsymbol{\beta} + \displaystyle\sum_{k \in \mathcal{K}} \boldsymbol{b}^\top \boldsymbol{\beta}^k \\
\text{subject to} & \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} \\
& \overline{\boldsymbol{\gamma}}^k \in \mathbb{R}^{N_\xi},\ \overline{\boldsymbol{x}}^k \in \mathbb{R}^{N_x}_+,\ \overline{\boldsymbol{w}}^k \in \mathbb{R}^{N_\xi}_+,\ \overline{\boldsymbol{y}}^k \in \mathbb{R}^{N_y}_+,\ k \in \mathcal{K} \\
& \mathbf{e}^\top \boldsymbol{\alpha} = 1,\ \boldsymbol{A}^\top \boldsymbol{\beta} = \displaystyle\sum_{k \in \mathcal{K}} \overline{\boldsymbol{\gamma}}^k \\
& \boldsymbol{A}^\top \boldsymbol{\beta}^k + \overline{\boldsymbol{\gamma}}^k = \boldsymbol{C} \overline{\boldsymbol{x}}^k + \boldsymbol{D} \overline{\boldsymbol{w}}^k + \boldsymbol{Q} \overline{\boldsymbol{y}}^k \quad \forall 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} \\
& \left.\begin{array}{l}
\overline{\boldsymbol{x}}^k \le \boldsymbol{x},\ \overline{\boldsymbol{x}}^k \le \boldsymbol{\alpha}_k \mathbf{e},\ \overline{\boldsymbol{x}}^k \ge (\boldsymbol{\alpha}_k - 1)\mathbf{e} + \boldsymbol{x} \\
\overline{\boldsymbol{w}}^k \le \boldsymbol{w},\ \overline{\boldsymbol{w}}^k \le \boldsymbol{\alpha}_k \mathbf{e},\ \overline{\boldsymbol{w}}^k \ge (\boldsymbol{\alpha}_k - 1)\mathbf{e} + \boldsymbol{w} \\
\overline{\boldsymbol{y}}^k \le \boldsymbol{y}^k,\ \overline{\boldsymbol{y}}^k \le \boldsymbol{\alpha}_k \mathbf{e},\ \overline{\boldsymbol{y}}^k \ge (\boldsymbol{\alpha}_k - 1)\mathbf{e} + \boldsymbol{y}^k \\
\overline{\boldsymbol{\gamma}}^k \le \boldsymbol{\gamma}^k + M(\mathbf{e} - \boldsymbol{w}),\ \overline{\boldsymbol{\gamma}}^k \le M\boldsymbol{w},\ \overline{\boldsymbol{\gamma}}^k \ge -M\boldsymbol{w},\ \overline{\boldsymbol{\gamma}}^k \ge \boldsymbol{\gamma}^k - M(\mathbf{e} - \boldsymbol{w})
\end{array}\right\} \quad \forall k \in \mathcal{K},
\end{array}
\tag{11}
$$

where M is a suitably chosen "big-M" constant.

We emphasize that the size of the MBLP (11) in Theorem 3 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), 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.

Remark 3. *Most MILP solvers*[4] *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. We leverage some of these computational tools in our experiments, see Section 10.*

Remark 4. *Suppose that we are only in the presence of exogenous uncertainty, i.e., $\boldsymbol{w} = \mathbf{e}$, $\boldsymbol{D} = \mathbf{0}$, and $\boldsymbol{V} = \mathbf{0}$. Then, the last set of constraints in Problem (11) implies that $\overline{\boldsymbol{\gamma}}^k = \boldsymbol{\gamma}^k$ for all $k \in \mathcal{K}$. Since $\boldsymbol{\gamma}^k$ is free, the second and third constraints are equivalent to*

$$\boldsymbol{A}^\top \boldsymbol{\beta} = \sum_{k \in \mathcal{K}} \boldsymbol{C} \overline{\boldsymbol{x}}^k + \boldsymbol{Q} \overline{\boldsymbol{y}}^k - \boldsymbol{A}^\top \boldsymbol{\beta}^k.$$

Exploiting the fact that $\boldsymbol{\alpha} \in \mathbb{R}^K_+$, $\mathbf{e}^\top \boldsymbol{\alpha} = 1$, and $\overline{\boldsymbol{x}}^k = \boldsymbol{\alpha}_k \boldsymbol{x}$, we can equivalently express this constraint as

$$\boldsymbol{A}^\top \left(\boldsymbol{\beta} + \sum_{k \in \mathcal{K}} \boldsymbol{\beta}^k \right) = \boldsymbol{C}\boldsymbol{x} + \sum_{k \in \mathcal{K}} \boldsymbol{Q} \overline{\boldsymbol{y}}^k.$$

We conclude that, in the presence of only exogenous uncertainty, Problem (11) *is equivalent to*

$$\begin{aligned} \text{minimize} \quad & \boldsymbol{b}^\top \boldsymbol{\beta} \\ \text{subject to} \quad & \boldsymbol{x} \in \mathcal{X},\ \boldsymbol{w} \in \mathcal{W},\ \boldsymbol{y}^k \in \mathcal{Y},\ k \in \mathcal{K} \\ & \boldsymbol{\alpha} \in \mathbb{R}_+^K,\ \boldsymbol{\beta} \in \mathbb{R}_+^R,\ \boldsymbol{\beta}^k \in \mathbb{R}_+^R,\ \overline{\boldsymbol{y}}^k \in \mathbb{R}_+^{N_y},\ k \in \mathcal{K} \\ & \mathbf{e}^\top \boldsymbol{\alpha} = 1,\ \boldsymbol{A}^\top \boldsymbol{\beta} = \boldsymbol{C}\boldsymbol{x} + \sum_{k \in \mathcal{K}} \boldsymbol{Q}\overline{\boldsymbol{y}}^k \\ & \boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{h} \quad \forall k \in \mathcal{K} \\ & \overline{\boldsymbol{y}}^k \leq \boldsymbol{y}^k,\ \overline{\boldsymbol{y}}^k \leq \boldsymbol{\alpha}_k \mathbf{e},\ \overline{\boldsymbol{y}}^k \geq (\boldsymbol{\alpha}_k - 1)\mathbf{e} + \boldsymbol{y}^k \quad \forall k \in \mathcal{K}, \end{aligned}$$

where we used the change of variables $\boldsymbol{\beta} \leftarrow \boldsymbol{\beta} + \sum_{k \in \mathcal{K}} \boldsymbol{\beta}^k$*. We then recover the MBLP formulation of the* K*-adaptability problem* (6)*. Thus, our reformulation encompasses as a special case the one from Hanasusanto et al. (2015).*

5. The K-Adaptability Problem with Constraint Uncertainty

The starting point of our analysis is the reformulation of the K-adaptability Problem ($\mathcal{P}_K$) as the min-max-min problem (7). Unfortunately, this problem is generally hard as testified by the following theorem.

Theorem 4. *Evaluating the objective function in Problem* (7) *if* K *is not fixed is strongly NP-hard.*

We reformulate Problem (7) equivalently by shifting the second-stage constraints $\boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{H}\boldsymbol{\xi}^k$ 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* (7)*, is equivalent to*

$$\begin{aligned} \text{minimize} \quad & \max_{\boldsymbol{\ell} \in \mathcal{L}} \ \max_{\{\boldsymbol{\xi}^k\}_{k \in \mathcal{K}} \in \Xi^K(\boldsymbol{w}, \boldsymbol{\ell})} \ \min_{\substack{k \in \mathcal{K}: \\ \boldsymbol{\ell}_k = 0}} \left\{ (\boldsymbol{\xi}^k)^\top \boldsymbol{C}\, \boldsymbol{x} + (\boldsymbol{\xi}^k)^\top \boldsymbol{D}\, \boldsymbol{w} + (\boldsymbol{\xi}^k)^\top \boldsymbol{Q}\, \boldsymbol{y}^k \right\} \\ \text{subject to} \quad & \boldsymbol{x} \in \mathcal{X},\ \boldsymbol{w} \in \mathcal{W},\ \boldsymbol{y}^k \in \mathcal{Y},\ k \in \mathcal{K}, \end{aligned} \tag{12}$$

where $\mathcal{L} := \{1, \ldots, L\}^K$*,* 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*

$$\Xi^K(\boldsymbol{w}, \boldsymbol{\ell}) := \left\{ \{\boldsymbol{\xi}^k\}_{k \in \mathcal{K}} \in \Xi^K : \begin{array}{ll} \boldsymbol{w} \circ \boldsymbol{\xi}^k = \boldsymbol{w} \circ \overline{\boldsymbol{\xi}} & \forall k \in \mathcal{K} \text{ for some } \overline{\boldsymbol{\xi}} \in \Xi \\ \boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{H}\boldsymbol{\xi}^k & \forall k \in \mathcal{K} : \boldsymbol{\ell}_k = 0 \\ [\boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k]_{\boldsymbol{\ell}_k} > [\boldsymbol{H}\boldsymbol{\xi}^k]_{\boldsymbol{\ell}_k} & \forall k \in \mathcal{K} : \boldsymbol{\ell}_k \neq 0 \end{array} \right\},$$

where, for notational 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 $\overline{\boldsymbol{\xi}}$. Thus, policy $\boldsymbol{y}^k$ is feasible in Problem (7) (and thus in Problem $(\mathcal{P}_K)$) if $\boldsymbol{\ell}_k = 0$. On the other hand, policy $\boldsymbol{y}^k$ violates the $\boldsymbol{\ell}_k$-th constraint in Problem (7) if $\boldsymbol{\ell}_k \neq 0$. Thus, if $\boldsymbol{\ell}_k \neq 0$, this implies that the $\boldsymbol{\ell}_k$-th constraint in $(\mathcal{P}_K)$ is violated for some $\boldsymbol{\xi} \in \Xi(\boldsymbol{w},\overline{\boldsymbol{\xi}})$ 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), $\boldsymbol{\ell}_k = 0$ if and only if policy $\boldsymbol{y}^k$ is *robustly* feasible in $(\mathcal{P}_K)$.

Remark 5. *We remark that instead of interchanging the inner minimization and maximization problems in $(\mathcal{P}_K)$, as is done in Lemma 1, we could robustify the constraints in the inner maximization problem in $(\mathcal{P}_K)$ to obtain a min-max-min-max problem where the inner maximization problem involves only a finite number of constraints parameterized by $\overline{\boldsymbol{\xi}}$. The resulting inner maximization problem involves products of $\overline{\boldsymbol{\xi}}$ with the dual variables resulting from the robustification. Interchanging the inner minimization and maximization problems then yields a non-convex bilinear maximization problem which precludes the use of standard robust optimization techniques. Similarly, dualizing the resulting inner maximization problem also results in a nonlinear non-convex formulation. Thus, we choose to first apply Lemma 1.*

Having brought Problem $(\mathcal{P}_K)$ to the form (12), it now presents a similar structure to a problem with objective uncertainty (see Section 4) 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 $\epsilon > 0$:

$$
\begin{array}{ll}
\text{minimize} & \displaystyle \max_{\boldsymbol{\ell}\in\mathcal{L}} \; \max_{\{\boldsymbol{\xi}^k\}_{k\in\mathcal{K}} \in \Xi^K_\epsilon(\boldsymbol{w},\boldsymbol{\ell})} \; \min_{\substack{k\in\mathcal{K}:\\ \boldsymbol{\ell}_k=0}} \left\{ (\boldsymbol{\xi}^k)^\top \boldsymbol{C}\,\boldsymbol{x} + (\boldsymbol{\xi}^k)^\top \boldsymbol{D}\,\boldsymbol{w} + (\boldsymbol{\xi}^k)^\top \boldsymbol{Q}\,\boldsymbol{y}^k \right\} \\
\text{subject to} & \boldsymbol{x}\in\mathcal{X},\; \boldsymbol{w}\in\mathcal{W},\; \boldsymbol{y}^k\in\mathcal{Y},\; k\in\mathcal{K},
\end{array}
\tag{12_ϵ}
$$

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

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

Using this definition, we next reformulate the approximate Problem (12_ϵ) equivalently as an MBLP.

Theorem 5. *The approximate problem* (12_ϵ) *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},\ \boldsymbol{\gamma}(\boldsymbol{\ell}) \in \mathbb{R}^K_+,\ \boldsymbol{\eta}^k(\boldsymbol{\ell}) \in \mathbb{R}^{N_\xi},\ k \in \mathcal{K},\ \boldsymbol{\ell} \in \mathcal{L} \\
& \left.\begin{array}{l}
\boldsymbol{\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}) = \displaystyle\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 \boldsymbol{\eta}^k(\boldsymbol{\ell}) = \boldsymbol{\lambda}_k(\boldsymbol{\ell})\,[\boldsymbol{C}\,\boldsymbol{x} + \boldsymbol{D}\,\boldsymbol{w} + \boldsymbol{Q}\,\boldsymbol{y}^k] \quad \forall k \in \mathcal{K} : \boldsymbol{\ell}_k = 0 \\
\boldsymbol{A}^\top \boldsymbol{\alpha}^k(\boldsymbol{\ell}) + [\boldsymbol{H}]_{\boldsymbol{\ell}_k} \gamma_k(\boldsymbol{\ell}) + \boldsymbol{w} \circ \boldsymbol{\eta}^k(\boldsymbol{\ell}) = \boldsymbol{\lambda}_k(\boldsymbol{\ell})\,[\boldsymbol{C}\,\boldsymbol{x} + \boldsymbol{D}\,\boldsymbol{w} + \boldsymbol{Q}\,\boldsymbol{y}^k] \quad \forall k \in \mathcal{K} : \boldsymbol{\ell}_k \neq 0 \\
\tau \geq \boldsymbol{b}^\top \left(\boldsymbol{\alpha}(\boldsymbol{\ell}) + \displaystyle\sum_{k \in \mathcal{K}} \boldsymbol{\alpha}^k(\boldsymbol{\ell}) \right) - \displaystyle\sum_{\substack{k \in \mathcal{K}: \\ \boldsymbol{\ell}_k = 0}} (\boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k)^\top \boldsymbol{\beta}^k(\boldsymbol{\ell}) \\
\qquad\qquad + \displaystyle\sum_{\substack{k \in \mathcal{K}: \\ \boldsymbol{\ell}_k \neq 0}} \left(\left[\boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \right]_{\boldsymbol{\ell}_k} - \epsilon \right) \gamma_k(\boldsymbol{\ell})
\end{array}\right\} \quad \forall \boldsymbol{\ell} \in \partial\mathcal{L} \\
& \left.\begin{array}{l}
\boldsymbol{A}^\top \boldsymbol{\alpha}(\boldsymbol{\ell}) = \displaystyle\sum_{k \in \mathcal{K}} \boldsymbol{w} \circ \boldsymbol{\eta}^k(\boldsymbol{\ell}) \\
\boldsymbol{A}^\top \boldsymbol{\alpha}^k(\boldsymbol{\ell}) + [\boldsymbol{H}]_{\boldsymbol{\ell}_k(\boldsymbol{\ell})} \gamma_k(\boldsymbol{\ell}) + \boldsymbol{w} \circ \boldsymbol{\eta}^k(\boldsymbol{\ell}) = \boldsymbol{0} \quad \forall k \in \mathcal{K} \\
\boldsymbol{b}^\top \left(\boldsymbol{\alpha}(\boldsymbol{\ell}) + \displaystyle\sum_{k \in \mathcal{K}} \boldsymbol{\alpha}^k(\boldsymbol{\ell}) \right) + \displaystyle\sum_{k \in \mathcal{K}} \left(\left[\boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \right]_{\boldsymbol{\ell}_k} - \epsilon \right) \gamma_k(\boldsymbol{\ell}) \leq -1
\end{array}\right\} \quad \forall \boldsymbol{\ell} \in \mathcal{L}_+,
\end{array}
\tag{13}
$$

where $\Lambda_K(\boldsymbol{\ell}) := \{\boldsymbol{\lambda} \in \mathbb{R}^K_+ \,:\, \mathbf{e}^\top \boldsymbol{\lambda} = 1,\ \boldsymbol{\lambda}_k = 0\ \forall k \in \mathcal{K} : \boldsymbol{\ell}_k \neq 0\}$*,* $\partial\mathcal{L} := \{\boldsymbol{\ell} \in \mathcal{L} : \boldsymbol{\ell} \ngtr \boldsymbol{0}\}$ *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* (12)*, respectively.*

Since all bilinear terms in Problem (13) involve one continuous and one binary variable, the problem can be reformulated equivalently as an MBLP using standard big-M techniques, see Hillier (2012). Similar to the robust counterpart resulting from the decision rule approximation proposed in Vayanos et al. (2011) (see also Section 3.3), Problem (13) 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 6. *Theorem 5 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* $\boldsymbol{H}$*. Suppose for example that* $\boldsymbol{T}(\boldsymbol{\xi}) := \sum_{n=1}^{N_\xi} \boldsymbol{T}^n \boldsymbol{\xi}_n$*,* $\boldsymbol{V}(\boldsymbol{\xi}) := \sum_{n=1}^{N_\xi} \boldsymbol{V}^n \boldsymbol{\xi}_n$ *and* $\boldsymbol{W}(\boldsymbol{\xi}) := \sum_{n=1}^{N_\xi} \boldsymbol{W}^n \boldsymbol{\xi}_n$ *for some* $\boldsymbol{T}^n \in \mathbb{R}^{L\times N_x}$*,* $\boldsymbol{V}^n \in \mathbb{R}^{L\times N_w}$ *and* $\boldsymbol{W}^n \in \mathbb{R}^{L\times N_y}$*,* $n = 1,\ldots,N_\xi$*. Using similar arguments as in the proof of Theorem 5, we can derive a formulation akin to the one in Problem* (13) *where* $\boldsymbol{T} = \boldsymbol{0}$*,* $\boldsymbol{V} = \boldsymbol{0}$*, and* $\boldsymbol{W} = \boldsymbol{0}$ *and, in the constraint associated with* $k \in \mathcal{K}$*,* $\boldsymbol{H}$ *is replaced with*

$$\boldsymbol{H} - [\boldsymbol{T}^1\boldsymbol{x}\cdots\boldsymbol{T}^{N_\xi}\boldsymbol{x}] - [\boldsymbol{V}^1\boldsymbol{w}\cdots\boldsymbol{V}^{N_\xi}\boldsymbol{w}] - [\boldsymbol{W}^1\boldsymbol{y}^k\cdots\boldsymbol{W}^{N_\xi}\boldsymbol{y}^k].$$

The following observation exactly parallels Remark 4 for the case of constraint uncertainty.

Observation 2. *Suppose that we are only in the presence of exogenous uncertainty, i.e.,* $\boldsymbol{w} = \mathbf{e}$*,* $\boldsymbol{D} = \boldsymbol{0}$*, and* $\boldsymbol{V} = \boldsymbol{0}$*. Then, Problem* (13) *reduces to the MBLP formulation of the K-adaptability problem* (6) *from Hanasusanto et al. (2015).*

6. 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, see Section 6.1. We generalize the K-adaptability approximation approach to this problem class by obtaining an equivalent reformulation of the K-adaptabiliy counterpart in the form of an MBLP in Section 6.2. As the size of this MBLP is exponential in K, we propose an efficient column-and-constraint generation procedure to address it in Section 6.3. We show that this formulation can be leveraged to solve certain classes of worst-case absolute regret minimization problems in Section 6.4.

6.1. 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

$$\max_{i\in\mathcal{I}} \; \boldsymbol{\xi}^\top \boldsymbol{C}^i\, \boldsymbol{x} + \boldsymbol{\xi}^\top \boldsymbol{D}^i\, \boldsymbol{w} + \boldsymbol{\xi}^\top \boldsymbol{Q}^i\, \boldsymbol{y}, \tag{14}$$

for some matrices $\boldsymbol{C}^i \in \mathbb{R}^{N_\xi \times N_x}$, $\boldsymbol{D}^i \in \mathbb{R}^{N_\xi \times N_\xi}$, 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, convex piecewise linear objective given by (14), and objective uncertainty is then expressible as

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

Note that, as in Section 4, our framework remains applicable in the presence of joint deterministic constraints $\boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y} \leq \boldsymbol{h}$ on the first and second stage decision variables. We omit these to minimize notational overhead.

6.2. K-Adaptability Approximation & MBLP Reformulation

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

$$\begin{array}{ll} \min & \max\limits_{\overline{\boldsymbol{\xi}} \in \Xi} \; \min\limits_{k \in \mathcal{K}} \; \max\limits_{\boldsymbol{\xi} \in \Xi(\boldsymbol{w}, \overline{\boldsymbol{\xi}})} \; \left\{ \max\limits_{i \in \mathcal{I}} \; \boldsymbol{\xi}^\top \boldsymbol{C}^i \, \boldsymbol{x} + \boldsymbol{\xi}^\top \boldsymbol{D}^i \, \boldsymbol{w} + \boldsymbol{\xi}^\top \boldsymbol{Q}^i \, \boldsymbol{y}^k \right\} \\ \text{s.t.} & \boldsymbol{x} \in \mathcal{X}, \; \boldsymbol{w} \in \mathcal{W}, \; \boldsymbol{y}^k \in \mathcal{Y}, \; k \in \mathcal{K}. \end{array} \tag{$\mathcal{PO}_K^{\rm PWL}$}$$

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 ($\mathcal{PO}_K^{\rm PWL}$), provided we index $\boldsymbol{\xi}$ by k.

Lemma 3. *The K-adaptability counterpart of Problem ($\mathcal{PO}_K^{\rm PWL}$) is equivalent to*

$$\begin{array}{ll} \text{minimize} & \max\limits_{\{\boldsymbol{\xi}^k\}_{k \in \mathcal{K}} \in \Xi^K(\boldsymbol{w})} \; \min\limits_{k \in \mathcal{K}} \; \left\{ \max\limits_{i \in \mathcal{I}} \; (\boldsymbol{\xi}^k)^\top \boldsymbol{C}^i \, \boldsymbol{x} + (\boldsymbol{\xi}^k)^\top \boldsymbol{D}^i \, \boldsymbol{w} + (\boldsymbol{\xi}^k)^\top \boldsymbol{Q}^i \, \boldsymbol{y}^k \right\} \\ \text{subject to} & \boldsymbol{x} \in \mathcal{X}, \; \boldsymbol{w} \in \mathcal{W}, \; \boldsymbol{y}^k \in \mathcal{Y}, \; k \in \mathcal{K}. \end{array} \tag{15}$$

Next, by leveraging Lemma 3, we are able to reformulate Problem (15) exactly as an MBLP. This result is summarized in the following theorem.

Theorem 6. *Problem ($\mathcal{PO}_K^{\rm PWL}$) is equivalent to the bilinear program*

$$\begin{array}{ll} \text{minimize} & \tau \\ \text{subject to} & \tau \in \mathbb{R}, \; \boldsymbol{x} \in \mathcal{X}, \; \boldsymbol{w} \in \mathcal{W}, \; \boldsymbol{y}^k \in \mathcal{Y}, \; k \in \mathcal{K} \\ & \boldsymbol{\alpha}^{\boldsymbol{i}} \in \mathbb{R}^K_+, \; \boldsymbol{\beta}^{\boldsymbol{i}} \in \mathbb{R}^R_+, \; \boldsymbol{\beta}^{\boldsymbol{i},k} \in \mathbb{R}^R_+, \; \boldsymbol{\gamma}^{\boldsymbol{i},k} \in \mathbb{R}^{N_\xi}, \; \forall k \in \mathcal{K}, \; \boldsymbol{i} \in \mathcal{I}^K \\ & \left. \begin{array}{l} \tau \geq \boldsymbol{b}^\top \boldsymbol{\beta}^{\boldsymbol{i}} + \sum\limits_{k \in \mathcal{K}} \boldsymbol{b}^\top \boldsymbol{\beta}^{\boldsymbol{i},k} \\ \mathrm{e}^\top \boldsymbol{\alpha}^{\boldsymbol{i}} = 1 \\ \boldsymbol{A}^\top \boldsymbol{\beta}^{\boldsymbol{i},k} + \boldsymbol{w} \circ \boldsymbol{\gamma}^{\boldsymbol{i},k} = \boldsymbol{\alpha}^{\boldsymbol{i}}_k \left(\boldsymbol{C}^{\boldsymbol{i}_k} \boldsymbol{x} + \boldsymbol{D}^{\boldsymbol{i}_k} \boldsymbol{w} + \boldsymbol{Q}^{\boldsymbol{i}_k} \boldsymbol{y}^k \right) \quad \forall k \in \mathcal{K} \\ \boldsymbol{A}^\top \boldsymbol{\beta}^{\boldsymbol{i}} = \sum\limits_{k \in \mathcal{K}} \boldsymbol{w} \circ \boldsymbol{\gamma}^{\boldsymbol{i},k} \end{array} \right\} \; \forall \boldsymbol{i} \in \mathcal{I}^K, \end{array} \tag{16}$$

which can be written equivalently as an MBLP by linearizing the bilinear terms, provided $\mathcal{X} \subseteq \{0,1\}^{N_x}$.

Albeit Problem (16) 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 ($K \gtrapprox 4$). In the remainder of this section, we exploit the specific structure of Problem ($\mathcal{PO}^{\rm PWL}$) to solve its K-adaptability counterpart *exactly* by reformulating it as an MBLP that presents an attractive structure amenable to decomposition techniques.

6.3. "Column-and-Constraint Generation" Algorithm

Recently, a so-called "column-and-constraint generation" algorithm has been proposed by Zeng and Zhao (2013) to solve two-stage linear robust optimization problems exactly. Here, we propose a variant of their algorithm to solve the K-adaptability counterpart ($\mathcal{PO}_K^{\rm PWL}$). 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 now detail this approach.

We define the following relaxed master problem parameterized by the index set $\widetilde{\mathcal{I}} \subseteq \mathcal{I}^K$

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

Given variables $(\tau, \boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}^k\}_{k\in\mathcal{K}})$ feasible in the master problem, we define the $\boldsymbol{i}$th subproblem, $\boldsymbol{i} \in \mathcal{I}$, through

$$
\begin{array}{ll}
\text{minimize} & 0 \\
\text{subject to} & \boldsymbol{\alpha}^{\boldsymbol{i}} \in \mathbb{R}^K_+,\ \boldsymbol{\beta}^{\boldsymbol{i}} \in \mathbb{R}^R_+,\ \boldsymbol{\beta}^{\boldsymbol{i},k} \in \mathbb{R}^R_+,\ \boldsymbol{\gamma}^{\boldsymbol{i},k} \in \mathbb{R}^{N_\xi},\ \forall k \in \mathcal{K} \\
& \tau \geq \boldsymbol{b}^\top \boldsymbol{\beta}^{\boldsymbol{i}} + \sum_{k\in\mathcal{K}} \boldsymbol{b}^\top \boldsymbol{\beta}^{\boldsymbol{i},k} \\
& \mathbf{e}^\top \boldsymbol{\alpha}^{\boldsymbol{i}} = 1 \\
& \boldsymbol{A}^\top \boldsymbol{\beta}^{\boldsymbol{i},k} + \boldsymbol{w} \circ \boldsymbol{\gamma}^{\boldsymbol{i},k} = \boldsymbol{\alpha}^{\boldsymbol{i}}_k \left(\boldsymbol{C}^{\boldsymbol{i}_k} \boldsymbol{x} + \boldsymbol{D}^{\boldsymbol{i}_k} \boldsymbol{w} + \boldsymbol{Q}^{\boldsymbol{i}_k} \boldsymbol{y}^k \right) \quad \forall k \in \mathcal{K} \\
& \boldsymbol{A}^\top \boldsymbol{\beta}^{\boldsymbol{i}} = \sum_{k\in\mathcal{K}} \boldsymbol{w} \circ \boldsymbol{\gamma}^{\boldsymbol{i},k}.
\end{array}
\qquad (\mathcal{CCG}^{\boldsymbol{i}}_{\text{sub}}(\tau, \boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}^k\}_{k\in\mathcal{K}}))
$$

An inspection of the Proof of Theorem 6 reveals that the last three constraints in Problem $(\mathcal{CCG}^{\boldsymbol{i}}_{\text{sub}}(\tau, \boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}^k\}_{k\in\mathcal{K}}))$ define the feasible set of the dual of a linear program that is feasible and bounded. Thus, for τ sufficiently large, Problem $(\mathcal{CCG}^{\boldsymbol{i}}_{\text{sub}}(\tau, \boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}^k\}_{k\in\mathcal{K}}))$ will be feasible.

To identify indices of subproblems $(\mathcal{CCG}^{\boldsymbol{i}}_{\text{sub}}(\tau, \boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}^k\}_{k\in\mathcal{K}}))$ that, given a solution $(\tau, \boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}^k\}_{k\in\mathcal{K}})$ to the relaxed master problem, are infeasible, we solve a *single* feasibility MBLP defined through

$$
\begin{array}{ll}
\max & \theta \\
\text{s.t.} & \theta \in \mathbb{R},\ \overline{\boldsymbol{\xi}} \in \Xi,\ \boldsymbol{\xi}^k \in \Xi(\boldsymbol{w}, \overline{\boldsymbol{\xi}}),\ k \in \mathcal{K} \\
& \boldsymbol{\eta} \in \mathbb{R}^K,\ \boldsymbol{\zeta}^k \in \{0,1\}^I,\ k \in \mathcal{K} \\
& \theta \leq \boldsymbol{\eta}_k \quad \forall k \in \mathcal{K} \\
& \left.\begin{array}{l}
\boldsymbol{\eta}_k \geq (\boldsymbol{\xi}^k)^\top \boldsymbol{C}^i \boldsymbol{x} + (\boldsymbol{\xi}^k)^\top \boldsymbol{D}^i \boldsymbol{w} + (\boldsymbol{\xi}^k)^\top \boldsymbol{Q}^i \boldsymbol{y}^k \\
\boldsymbol{\eta}_k \leq (\boldsymbol{\xi}^k)^\top \boldsymbol{C}^i \boldsymbol{x} + (\boldsymbol{\xi}^k)^\top \boldsymbol{D}^i \boldsymbol{w} + (\boldsymbol{\xi}^k)^\top \boldsymbol{Q}^i \boldsymbol{y}^k + M(1 - \boldsymbol{\zeta}^k_i)
\end{array}\right\} \begin{array}{l} \forall i \in \mathcal{I}, \\ k \in \mathcal{K} \end{array} \\
& \mathbf{e}^\top \boldsymbol{\zeta}^k = 1 \quad \forall k \in \mathcal{K}.
\end{array}
\qquad (\mathcal{CCG}_{\text{feas}}(\boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}^k\}_{k\in\mathcal{K}}))
$$

The following proposition enables us to bound the optimality gap associated with a given feasible solution to the relaxed master problem.

Proposition 2. *Let $(\boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}^k\}_{k\in\mathcal{K}})$ be feasible in the relaxed master problem $(\mathcal{CCG}_{\text{mstr}}(\widetilde{\mathcal{I}}))$. Then, $(\boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}^k\}_{k\in\mathcal{K}})$ is feasible in Problem $(\mathcal{PO}^{\text{PWL}}_K)$ and the objective value of $(\boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}^k\}_{k\in\mathcal{K}})$ in Problem $(\mathcal{PO}^{\text{PWL}}_K)$ is given by the optimal objective value of Problem $(\mathcal{CCG}_{\text{feas}}(\boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}^k\}_{k\in\mathcal{K}}))$.*

Proposition 2 implies that, for any $(\boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}^k\}_{k\in\mathcal{K}})$ feasible in the relaxed master problem $(\mathcal{CCG}_{\text{mstr}}(\widetilde{\mathcal{I}}))$, the optimal value of $(\mathcal{CCG}_{\text{feas}}(\boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}^k\}_{k\in\mathcal{K}}))$ yields an upper bound to the optimal value of the K-adaptability problem ($\mathcal{PO}_K^{\mathrm{PWL}}$). At the same time, it is evident that for any index set $\widetilde{\mathcal{I}} \subseteq \mathcal{I}^K$, the optimal value of Problem ($\mathcal{CCG}_{\mathrm{mstr}}(\widetilde{\mathcal{I}})$) yields a lower bound to the optimal objective value of Problem ($\mathcal{PO}_K^{\mathrm{PWL}}$).

The lemma below is key to identify indices of subproblems $\boldsymbol{i} \in \mathcal{I}^K$ that are infeasible.

Lemma 4. *Let* $(\tau, \boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}^k\}_{k\in\mathcal{K}}, \{\boldsymbol{\alpha}^{\boldsymbol{i}}, \boldsymbol{\beta}^{\boldsymbol{i}}\}_{\boldsymbol{i}\in\widetilde{\mathcal{I}}}, \{\boldsymbol{\beta}^{\boldsymbol{i},k}, \boldsymbol{\gamma}^{\boldsymbol{i},k}\}_{\boldsymbol{i}\in\widetilde{\mathcal{I}},k\in\mathcal{K}})$ *be optimal in the relaxed master problem* ($\mathcal{CCG}_{\mathrm{mstr}}(\widetilde{\mathcal{I}})$). *Let* $(\theta, \overline{\boldsymbol{\xi}}, \{\boldsymbol{\xi}^k\}_{k\in\mathcal{K}}, \boldsymbol{\eta}, \{\boldsymbol{\zeta}^k\}_{k\in\mathcal{K}})$ *be optimal in Problem* ($\mathcal{CCG}_{\mathrm{feas}}(\boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}^k\}_{k\in\mathcal{K}})$). *Then, the following hold:*

(i) $\theta \geq \tau$;

(ii) If $\theta = \tau$, *then Problem* ($\mathcal{CCG}_{\mathrm{sub}}^{\boldsymbol{i}}(\tau, \boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}^k\}_{k\in\mathcal{K}})$) *is feasible for all* $\boldsymbol{i} \in \mathcal{I}^K$;

(iii) If $\theta > \tau$, *then the index* $\boldsymbol{i}$ *defined through*

$$\boldsymbol{i}_k := \sum_{i\in\mathcal{I}} i \cdot \mathbb{I}(\boldsymbol{\zeta}_i^k = 1) \quad \forall k \in \mathcal{K}$$

corresponds to an infeasible subproblem, i.e., Problem ($\mathcal{CCG}_{\mathrm{sub}}^{\boldsymbol{i}}(\tau, \boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}^k\}_{k\in\mathcal{K}})$) *is infeasible.*

Propositions 2 and Lemma 4 culminate in Algorithm 1 whose convergence is guaranteed by the following theorem.

Theorem 7. *Algorithm 1 terminates in a final number of steps with a feasible solution to Problem* ($\mathcal{PO}_K^{\mathrm{PWL}}$). *The objective value* θ *attained by this solution is within* δ *of the optimal objective value of the problem.*

In the following, we 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 ($\mathcal{PO}^{\mathrm{PWL}}$).

6.4. Application to Worst-Case Absolute Regret Minimization

According to the "worst-case absolute regret" criterion, the performance of a decision is evaluated with respect to the worst-case regret that is experienced, when comparing the performance of the decision taken relative to the performance of the best decision that should have been taken *in hindsight*, after all uncertain parameters are revealed, see e.g., Savage (1951). The minimization of worst-case absolute regret is often believed to mitigate the conservatism of classical robust optimization and is thus attractive in practical applications, see also Section 10 for corroborating evidence.

Mathematically, we are given a utility function

$$u(\boldsymbol{x}, \boldsymbol{w}, \boldsymbol{y}, \boldsymbol{\xi}) := \boldsymbol{\xi}^\top \boldsymbol{C}\, \boldsymbol{x} + \boldsymbol{\xi}^\top \boldsymbol{D}\, \boldsymbol{w} + \boldsymbol{\xi}^\top \boldsymbol{Q}\, \boldsymbol{y} \tag{17}$$

Algorithm 1: "Column-and-Constraint" Generation Procedure.

Inputs: Optimality tolerance δ; K-adaptability parameter K;

Output: Near optimal solution $(\boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}^k\}_{k\in\mathcal{K}})$ to Problem $(\mathcal{PO}_K^{\mathrm{PWL}})$ with associated objective θ;

Initialization:

Initialize upper and lower bounds: $\mathrm{LB} \leftarrow -\infty$ and $\mathrm{UB} \leftarrow +\infty$;

Initialize index set: $\widetilde{\mathcal{I}} \leftarrow \{\mathbf{e}\}$;

while $\mathrm{UB} - \mathrm{LB} > \delta$ **do**

Solve the master problem $(\mathcal{CCG}_{\mathrm{mstr}}(\widetilde{\mathcal{I}}))$, let $(\tau, \boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}^k\}_{k\in\mathcal{K}}, \{\boldsymbol{\alpha}^{\boldsymbol{i}}, \boldsymbol{\beta}^{\boldsymbol{i}}\}_{\boldsymbol{i}\in\widetilde{\mathcal{I}}}, \{\boldsymbol{\beta}^{\boldsymbol{i},k}, \boldsymbol{\gamma}^{\boldsymbol{i},k}\}_{\boldsymbol{i}\in\widetilde{\mathcal{I}}, k\in\mathcal{K}})$ be an optimal solution;

Let $\mathrm{LB} \leftarrow \tau$;

Solve the feasibility subproblem $(\mathcal{CCG}_{\mathrm{feas}}(\boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}^k\}_{k\in\mathcal{K}}))$, let $(\theta, \overline{\boldsymbol{\xi}}, \{\boldsymbol{\xi}^k\}_{k\in\mathcal{K}}, \boldsymbol{\eta}, \{\boldsymbol{\zeta}^k\}_{k\in\mathcal{K}})$ denote an optimal solution;

Let $\mathrm{UB} \leftarrow \theta$;

if $\theta > \tau$ **then**

$\boldsymbol{i}_k \leftarrow \sum_{i\in\mathcal{I}} i \cdot \mathbb{I}(\boldsymbol{\zeta}_i^k = 1)$ for all $k \in \mathcal{K}$;

$\widetilde{\mathcal{I}} \leftarrow \widetilde{\mathcal{I}} \cup \{\boldsymbol{i}\}$;

end

end

Result: $(\boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}^k\}_{k\in\mathcal{K}})$ is near-optimal in $(\mathcal{PO}_K^{\mathrm{PWL}})$ with objective value θ.

for which high values are preferred. This function depends on both the decisions $\boldsymbol{x}$, $\boldsymbol{w}$, and $\boldsymbol{y}$, and on the uncertain parameters $\boldsymbol{\xi}$. Given a realization of $\boldsymbol{\xi}$, we can measure the *absolute regret* of a decision $(\boldsymbol{x}, \boldsymbol{w}, \boldsymbol{y})$ as the difference between the utility of the best decision in hindsight (i.e., after $\boldsymbol{\xi}$ becomes known) and the utility of the decision actually taken, i.e.,

$$\left\{ \max_{\boldsymbol{x}', \boldsymbol{w}', \boldsymbol{y}'} \;\; u(\boldsymbol{x}', \boldsymbol{w}', \boldsymbol{y}', \boldsymbol{\xi}) - u(\boldsymbol{x}, \boldsymbol{w}, \boldsymbol{y}, \boldsymbol{\xi}) \,:\, \boldsymbol{x}' \in \mathcal{X},\ \boldsymbol{w}' \in \mathcal{W},\ \boldsymbol{y}' \in \mathcal{Y} \right\}.$$

Regret averse decision-makers seek to minimize the worst-case (maximum) absolute regret

$$\max_{\boldsymbol{\xi}\in\Xi(\boldsymbol{w},\overline{\boldsymbol{\xi}})} \left\{ \max_{\boldsymbol{x}', \boldsymbol{w}', \boldsymbol{y}'} \;\; u(\boldsymbol{x}', \boldsymbol{w}', \boldsymbol{y}', \boldsymbol{\xi}) - u(\boldsymbol{x}, \boldsymbol{w}, \boldsymbol{y}, \boldsymbol{\xi}) \,:\, \boldsymbol{x}' \in \mathcal{X},\ \boldsymbol{w}' \in \mathcal{W},\ \boldsymbol{y}' \in \mathcal{Y} \right\}. \tag{18}$$

A two-stage robust optimization problem with DDID in which the decision-maker seeks to minimize his worst-case absolute regret is then expressible as

$$\begin{aligned} \min \quad & \max_{\overline{\boldsymbol{\xi}} \in \Xi} \; \min_{\boldsymbol{y} \in \mathcal{Y}} \; \max_{\boldsymbol{\xi} \in \Xi(\boldsymbol{w}, \overline{\boldsymbol{\xi}})} \left\{ \max_{\boldsymbol{x}', \boldsymbol{w}', \boldsymbol{y}'} \; u(\boldsymbol{x}', \boldsymbol{w}', \boldsymbol{y}', \boldsymbol{\xi}) - u(\boldsymbol{x}, \boldsymbol{w}, \boldsymbol{y}, \boldsymbol{\xi}) \, : \, \boldsymbol{x}' \in \mathcal{X}, \, \boldsymbol{w}' \in \mathcal{W}, \, \boldsymbol{y}' \in \mathcal{Y} \right\} \\ \text{s.t.} \quad & \boldsymbol{x} \in \mathcal{X}, \, \boldsymbol{w} \in \mathcal{W}. \end{aligned} \tag{$\mathcal{WCAR}$}$$

The following observation shows that under certain assumptions, Problem ($\mathcal{WCAR}$) can be written in the form ($\mathcal{PO}^{\text{PWL}}$).

Observation 3. *Suppose that the utility function u in* (17) *and the feasible sets $\mathcal{X}$, $\mathcal{W}$, and $\mathcal{Y}$ in Problem* ($\mathcal{WCAR}$) *are such that:*

(i) Either $\boldsymbol{C} = \mathbf{0}$ or $\mathcal{X} := \{\boldsymbol{x} : \mathbf{e}^\top \boldsymbol{x} = 1\}$, and

(ii) Either $\boldsymbol{D} = \mathbf{0}$ or $\mathcal{W} := \{\boldsymbol{w} : \mathbf{e}^\top \boldsymbol{w} = 1\}$, and

(iii) Either $\boldsymbol{Q} = \mathbf{0}$ or $\mathcal{Y} := \{\boldsymbol{y} : \mathbf{e}^\top \boldsymbol{y} = 1\}$.

Then, Problem ($\mathcal{WCAR}$) *can be written in the form* ($\mathcal{PO}^{\text{PWL}}$)*.*

In Sections 9 and 10, we leverage Observation 3 and use Theorems 6 and 7, and Algorithm 1 to solve an active preference learning problem that seeks to recommend housing allocation policies with least possible worst-case regret.

7. The Multi-Stage Case with Objective Uncertainty

In this section, we show that many of our results generalize to the multi-stage case. To this end, we propose a novel formulation of multi-stage robust optimization problems with decision-dependent information discovery. This formulation will underpin our ability to generalize the K-adaptability approximation approach for problems with endogenous uncertainty to the multi-stage setting. As in the two-stage case, it will enable us to construct more tractable and accurate conservative approximations to such problems. To minimize notational overhead, we focus on problems with only objective uncertainty. However, our results for the constraint uncertainty case generalize to this multi-stage setting too.

The remainder of this section is organized as follows. First, we introduce two equivalent models of multi-stage robust optimization problems with exogenous and endogenous uncertainty, see Sections 7.1 and 7.2, respectively. Then, in Section 7.3, we leverage the second formulation to generalize the K-adaptability framework to the multi-stage setting. Finally, we show in Section 7.4 that the K-adaptability problem can be reformulated equivalently as an MBLP.

7.1. Multi-Stage Robust Optimization with Exogenous Uncertainty

In the literature, and similar to the two-stage case, there are (broadly speaking) two formulations of a generic multi-stage robust optimization problem with exogenous uncertainty over the planning horizon $\mathcal{T} := \{1,\ldots,T\}$. These differ in the way in which the ability for the time t decisions to adapt to the history of observed parameter realizations is modeled.

Decision Rule Formulation. In the first model, one optimizes today over all recourse actions $\boldsymbol{y}^t(\boldsymbol{\xi}) \in \mathbb{R}^{N_{y_t}}$ that will be taken in each realization of $\boldsymbol{\xi} \in \Xi$. Under this modeling paradigm, a multi-stage robust optimization problem with *exogenous* uncertainty is expressible as

$$
\begin{array}{ll}
\text{minimize} & \displaystyle\max_{\boldsymbol{\xi}\in\Xi} \; \sum_{t\in\mathcal{T}} \boldsymbol{\xi}^\top \boldsymbol{Q}^t \, \boldsymbol{y}^t(\boldsymbol{\xi}) \\
\text{subject to} & \boldsymbol{y}^t \in \mathcal{L}_{N_\xi}^{N_{y_t}} \quad \forall t \in \mathcal{T} \\
& \left.\begin{array}{l}
\boldsymbol{y}^t(\boldsymbol{\xi}) \in \mathcal{Y}_t \quad \forall t \in \mathcal{T} \\
\displaystyle\sum_{t\in\mathcal{T}} \boldsymbol{W}^t \boldsymbol{y}^t(\boldsymbol{\xi}) \le \boldsymbol{H}\boldsymbol{\xi} \\
\boldsymbol{y}^t(\boldsymbol{\xi}) = \boldsymbol{y}^t(\boldsymbol{w}^{t-1} \circ \boldsymbol{\xi}) \quad \forall t \in \mathcal{T}
\end{array}\right\} \quad \forall \boldsymbol{\xi} \in \Xi,
\end{array}
\tag{19}
$$

where $\boldsymbol{Q}^t \in \mathbb{R}^{N_\xi \times N_{y_t}}$, $\boldsymbol{W}^t \in \mathbb{R}^{L_t \times N_{y_t}}$, and $\boldsymbol{H}^t \in \mathbb{R}^{L_t \times N_\xi}$. The *fixed* binary vector $\boldsymbol{w}^t \in \{0,1\}^{N_\xi}$ represents the *information base* at time $t+1$, i.e., it encodes the information revealed up to (and including) time t. Thus, $\boldsymbol{w}_i^t = 1$ if and only if $\boldsymbol{\xi}_i$ has been observed at some time $\tau \in \{0,\ldots,t\}$. As information cannot be forgotten, it holds that $\boldsymbol{w}^t \ge \boldsymbol{w}^{t-1}$ for all $t \in \mathcal{T}$. The last constraint in Problem (19) ensures that the decisions $\boldsymbol{y}^t$, $t \in \mathcal{T}$, are non-anticipative: it stipulates that $\boldsymbol{y}^t$ can only depend on those parameters that have been observed up to and including time $t-1$.

Dynamic Formulation. In the second model, the recourse decisions $\boldsymbol{y}^t$ are optimized explicitly *after* nature is done making a decision. Under this modeling paradigm, a generic multi-stage robust problem with exogenous uncertainty is expressible as:

$$
\begin{array}{l}
\displaystyle\min_{\boldsymbol{y}^1\in\mathcal{Y}_1} \max_{\boldsymbol{\xi}^1\in\Xi} \min_{\boldsymbol{y}^2\in\mathcal{Y}_2} \max_{\boldsymbol{\xi}^2\in\Xi(\boldsymbol{w}^1,\boldsymbol{\xi}^1)} \cdots \min_{\boldsymbol{y}^T\in\mathcal{Y}_T} \max_{\boldsymbol{\xi}^T\in\Xi(\boldsymbol{w}^{T-1},\boldsymbol{\xi}^{T-1})} \sum_{t\in\mathcal{T}} (\boldsymbol{\xi}^T)^\top \boldsymbol{Q}^t \, \boldsymbol{y}^t \\
\qquad\qquad\qquad\qquad\qquad \text{s.t.} \quad \displaystyle\sum_{t\in\mathcal{T}} \boldsymbol{W}^t \boldsymbol{y}^t \le \boldsymbol{H}(\boldsymbol{\xi}^T) \quad \forall \boldsymbol{\xi}^T \in \Xi(\boldsymbol{w}^{T-1},\boldsymbol{\xi}^{T-1}),
\end{array}
\tag{20}
$$

where, as in the two-stage case, we have

$$
\Xi(\boldsymbol{w}^{t-1},\boldsymbol{\xi}^{t-1}) := \left\{ \boldsymbol{\xi}^t \in \Xi \,:\, \boldsymbol{w}^{t-1} \circ \boldsymbol{\xi}^t = \boldsymbol{w}^{t-1} \circ \boldsymbol{\xi}^{t-1} \right\} \quad \forall t \in \mathcal{T}.
$$

We state the following theorem without proof. The proof follows directly from arguments similar to the proof of Theorem 1 by following a recursive argument. Thus, we omit it.

Theorem 8. *Problems* (19) *and* (20) *are equivalent.*

7.2. Multi-Stage Robust Optimization with Decision-Dependent Information Discovery

In this section, we investigate a variant of Problem (19) (and accordingly (20)) that enjoys much greater modeling flexibility since the time of information discovery (i.e., the information base) is kept flexible. Thus, we interpret the information base $\boldsymbol{w}^t \in \mathcal{W}_t \subseteq \{0,1\}^{N_\xi}$ as a decision variable, which is allowed to depend on $\boldsymbol{\xi}$. The set $\mathcal{W}_t$ may incorporate constraints stipulating, for example, that a specific uncertain parameter can only be observed after a certain stage or that an uncertain parameter can only be observed if another one has, etc. We assume that a cost is incurred for including uncertain parameters in the information base (equivalently, for observing uncertain parameters) and that the observation decisions $\boldsymbol{w}^t$ also impact the constraints through the additional term $\sum_{t\in\mathcal{T}} \boldsymbol{V}^t \boldsymbol{w}^t$, where $\boldsymbol{V}^t \in \mathbb{R}^{L_t \times N_\xi}$. As before, we propose two equivalent models for multi-stage robust problems with DDID which differ in the way the ability for the time t decisions to depend on the history of parameter realizations is modeled.

Decision Rule Formulation. In the first model, one optimizes today over all recourse actions $\boldsymbol{w}^t(\boldsymbol{\xi}) \in \mathbb{R}^{N_\xi}$ and $\boldsymbol{y}^t(\boldsymbol{\xi}) \in \mathbb{R}^{N_{y_t}}$ that will be taken in each realization of $\boldsymbol{\xi} \in \Xi$. Under this modeling paradigm, a multi-stage robust optimization problem with decision-dependent information discovery, originally proposed in Vayanos et al. (2011), reads

$$
\begin{array}{ll}
\text{minimize} & \displaystyle \max_{\boldsymbol{\xi}\in\Xi} \; \sum_{t\in\mathcal{T}} \boldsymbol{\xi}^\top \boldsymbol{D}^t\, \boldsymbol{w}^t(\boldsymbol{\xi}) + \boldsymbol{\xi}^\top \boldsymbol{Q}^t\, \boldsymbol{y}^t(\boldsymbol{\xi}) \\
\text{subject to} & \boldsymbol{w}^t \in \mathcal{L}^{N_\xi}_{N_\xi},\; \boldsymbol{y}^t \in \mathcal{L}^{N_{y_t}}_{N_\xi} \quad \forall t \in \mathcal{T} \\
& \left.\begin{array}{l}
\boldsymbol{w}^t(\boldsymbol{\xi}) \in \mathcal{W}_t,\; \boldsymbol{y}^t(\boldsymbol{\xi}) \in \mathcal{Y}_t \quad \forall t \in \mathcal{T} \\
\displaystyle\sum_{t\in\mathcal{T}} \boldsymbol{V}^t \boldsymbol{w}^t(\boldsymbol{\xi}) + \boldsymbol{W}^t \boldsymbol{y}^t(\boldsymbol{\xi}) \leq \boldsymbol{H}\boldsymbol{\xi} \\
\boldsymbol{w}^t(\boldsymbol{\xi}) \geq \boldsymbol{w}^{t-1}(\boldsymbol{\xi}) \quad \forall t \in \mathcal{T} \\
\boldsymbol{y}^t(\boldsymbol{\xi}) = \boldsymbol{y}^t(\boldsymbol{w}^{t-1}(\boldsymbol{\xi}) \circ \boldsymbol{\xi}) \quad \forall t \in \mathcal{T} \\
\boldsymbol{w}^t(\boldsymbol{\xi}) = \boldsymbol{w}^t(\boldsymbol{w}^{t-1}(\boldsymbol{\xi}) \circ \boldsymbol{\xi}) \quad \forall t \in \mathcal{T}
\end{array}\right\} \quad \forall \boldsymbol{\xi} \in \Xi,
\end{array}
\tag{21}
$$

where $\boldsymbol{w}^0(\boldsymbol{\xi}) = \boldsymbol{w}^0$ for all $\boldsymbol{\xi} \in \Xi$ and $\boldsymbol{w}^0$ is given and encodes the information available at the beginning of the planning horizon.

Dynamic Formulation. In the second model, the recourse decisions $\boldsymbol{w}^t$ and $\boldsymbol{y}^t$ are optimized explicitly *after* nature is done selecting the parameters we have chosen to observe in the past. Under this modeling paradigm, a generic multi-stage robust problem with DDID is expressible as:

$$\begin{aligned}\min_{\substack{\boldsymbol{y}^1\in\mathcal{Y}_1\\ \boldsymbol{w}^1\in\mathcal{W}_1}} \max_{\boldsymbol{\xi}^1\in\Xi} \min_{\substack{\boldsymbol{y}^2\in\mathcal{Y}_2\\ \boldsymbol{w}^2\in\mathcal{W}_2\\ \boldsymbol{w}^2\geq\boldsymbol{w}^1}} \max_{\boldsymbol{\xi}^2\in\Xi(\boldsymbol{w}^1,\boldsymbol{\xi}^1)} \min_{\substack{\boldsymbol{y}^3\in\mathcal{Y}_3\\ \boldsymbol{w}^3\in\mathcal{W}_3\\ \boldsymbol{w}^3\geq\boldsymbol{w}^2}} \cdots \min_{\boldsymbol{y}^T\in\mathcal{Y}_T} \max_{\boldsymbol{\xi}^T\in\Xi(\boldsymbol{w}^{T-1},\boldsymbol{\xi}^{T-1})} & \sum_{t\in\mathcal{T}} (\boldsymbol{\xi}^T)^\top \boldsymbol{D}^t\, \boldsymbol{w}^t + (\boldsymbol{\xi}^T)^\top \boldsymbol{Q}^t\, \boldsymbol{y}^t \\ \text{s.t.} \quad & \sum_{t\in\mathcal{T}} \boldsymbol{V}^t\boldsymbol{w}^t + \boldsymbol{W}^t\boldsymbol{y}^t \leq \boldsymbol{H}\boldsymbol{\xi}^T \quad \forall \boldsymbol{\xi}^T\in\Xi(\boldsymbol{w}^{T-1},\boldsymbol{\xi}^{T-1}).\end{aligned} \tag{$\mathcal{MP}$}$$

Similarly to the exogenous case, it can be shown that the two models above are equivalent as summarized in the following theorem.

Theorem 9. *Problems* (21) *and* ($\mathcal{MP}$) *are equivalent.*

The proof of Theorem 9 follows by applying Theorem 2 recursively and we thus omit it.

7.3. K-Adaptability for Multi-Stage Problems with Decision-Dependent Information Discovery

We henceforth propose to approximate Problem ($\mathcal{MP}$) with its K-adaptability counterpart, whereby K candidate policies are selected here-and-now (for each time period) and the best of these policies is selected, in an adaptive fashion, at each stage. To streamline presentation, we focus on the case where Problem ($\mathcal{MP}$) presents only objective uncertainty. Thus, the K-adaptability counterpart of the multi-stage robust problem ($\mathcal{MP}$) with DDID is expressible as

$$\begin{aligned}\min_{k_1\in\mathcal{K}} \max_{\boldsymbol{\xi}^1\in\Xi} \min_{k_2\in\mathcal{K}} \max_{\boldsymbol{\xi}^2\in\Xi(\boldsymbol{w}^{1,k_1},\boldsymbol{\xi}^1)} \cdots \min_{k_T\in\mathcal{K}} \max_{\boldsymbol{\xi}^T\in\Xi(\boldsymbol{w}^{T-1,k_{T-1}},\boldsymbol{\xi}^{T-1})} & \sum_{t\in\mathcal{T}} (\boldsymbol{\xi}^T)^\top \boldsymbol{D}^t\, \boldsymbol{w}^{t,k_t} + (\boldsymbol{\xi}^T)^\top \boldsymbol{Q}^t\, \boldsymbol{y}^{t,k_t} \\ \text{s.t.} \quad & \boldsymbol{y}^{t,k}\in\mathcal{Y}_t,\ \boldsymbol{w}^{t,k}\in\mathcal{W}_t \quad \forall t\in\mathcal{T},\ k\in\mathcal{K} \\ & \boldsymbol{w}^{t,k_t}\geq\boldsymbol{w}^{t-1,k_{t-1}} \quad \forall t\in\mathcal{T},\ k_t\in\mathcal{K},\ k_{t-1}\in\mathcal{K} \\ & \sum_{t\in\mathcal{T}} \boldsymbol{V}^t\boldsymbol{w}^{t,k_t} + \boldsymbol{W}^t\boldsymbol{y}^{t,k_t} \leq \boldsymbol{h} \quad \forall k_1,\ldots,k_T\in\mathcal{K},\end{aligned} \tag{$\mathcal{MPO}_K$}$$

where we have defined $\boldsymbol{w}^{0,k}=\boldsymbol{w}^0$ for all $k\in\mathcal{K}$ with $\boldsymbol{w}^0_i=1$ if and only if $\boldsymbol{\xi}_i$ is observed at the beginning of the planning horizon and, as in the two-stage case, we have moved the deterministic constraints to the first stage. While Problem ($\mathcal{MPO}_K$) appears significantly more complicated than its two-stage counterpart, it can be brought to a min-max-min form at the cost of lifting the dimension of the uncertainty, as shown in the following lemma.

Lemma 5. *The K-adaptability counterpart of the multi-stage robust optimization problem with decision-dependent information discovery, Problem ($\mathcal{MPO}_K$), is equivalent to the two-stage robust problem*

$$
\begin{array}{ll}
\text{minimize} & \displaystyle \max_{\substack{\boldsymbol{\xi}^{T,k_T\cdots k_1}\in\Xi^T(\boldsymbol{w}^{1,k_1},\ldots,\boldsymbol{w}^{T-1,k_{T-1}})\\ k_1,\ldots,k_T\in\mathcal{K}}} \; \min_{k_1,\ldots,k_T\in\mathcal{K}} \; \sum_{t\in\mathcal{T}} (\boldsymbol{\xi}^{T,k_T\cdots k_1})^\top \boldsymbol{D}^t\, \boldsymbol{w}^{t,k_t} + (\boldsymbol{\xi}^{T,k_T\cdots k_1})^\top \boldsymbol{Q}^t\, \boldsymbol{y}^{t,k_t} \\
\text{subject to} & \boldsymbol{y}^{t,k} \in \mathcal{Y}_t,\; \boldsymbol{w}^{t,k}\in\mathcal{W}_t \quad \forall t\in\mathcal{T},\, k\in\mathcal{K} \\
& \boldsymbol{w}^{t,k_t} \geq \boldsymbol{w}^{t-1,k_{t-1}} \quad \forall t\in\mathcal{T},\, k_t\in\mathcal{K},\, k_{t-1}\in\mathcal{K} \\
& \displaystyle\sum_{t\in\mathcal{T}} \boldsymbol{V}^t\boldsymbol{w}^{t,k_t} + \boldsymbol{W}^t\boldsymbol{y}^{t,k_t} \leq \boldsymbol{h} \quad \forall k_1,\ldots,k_T\in\mathcal{K},
\end{array}
\tag{22}
$$

where $\Xi^T(\boldsymbol{w}^1,\ldots,\boldsymbol{w}^{T-1}) := \{\boldsymbol{\xi}^T\in\Xi : \boldsymbol{\xi}^1\in\Xi,\, \boldsymbol{\xi}^t\in\Xi(\boldsymbol{w}^{t-1},\boldsymbol{\xi}^{t-1}) \;\; \forall t\in\mathcal{T}\backslash\{1\}\}$.

The proof of Lemma 5 follows directly by applying the proof of Lemma 1 iteratively, starting at the last period and inverting the order of the maximization and minimization by lifting the uncertain parameters to a higher dimensional space.

Observation 4. *For any fixed K and decision $(\boldsymbol{x}, \{\boldsymbol{w}^{t,k_t}\}_{t\in\mathcal{T},k_t\in\mathcal{K}}, \{\boldsymbol{y}^{t,k_t}\}_{t\in\mathcal{T},k_t\in\mathcal{K}})$, the optimal objective value of the K-adaptability problem* (22) *can be evaluated by solving an LP whose size is exponential in the size of the input; and in particular exponential in T.*

As the proof of Observation 4 exactly parallels that of Observation 1 for the two-stage case, we omit it.

7.4. Reformulation as a Mixed-Binary Linear Program

In Observation 4, we showed that for any fixed K, $\boldsymbol{x}$, $\{\boldsymbol{w}^{t,k_t}\}_{t\in\mathcal{T},k_t\in\mathcal{K}}$, and $\{\boldsymbol{y}^{t,k_t}\}_{t\in\mathcal{T},k_t\in\mathcal{K}}$, the objective function in Problem ($\mathcal{MPO}_K$) can be evaluated by means of a polynomially sized LP. By dualizing this LP and linearizing the resulting bilinear terms, we can obtain an equivalent reformulation of Problem ($\mathcal{MPO}_K$) in the form of a mixed-binary linear program (MBLP).

Theorem 10. *Suppose $\mathcal{Y}_t \subseteq \{0,1\}^{N_{y_t}}$ for all $t \in \mathcal{T}$. Then, Problem $(\mathcal{MPO}_K)$ is equivalent to the following bilinear program*

$$
\begin{array}{ll}
\text{minimize} & \displaystyle\sum_{t\in\mathcal{T}} \sum_{k_1\in\mathcal{K}} \cdots \sum_{k_t\in\mathcal{K}} \boldsymbol{b}^\top \boldsymbol{\gamma}^{t,k_1\cdots k_t} \\
\text{subject to} & \boldsymbol{\alpha} \in \mathbb{R}_+^{K^T},\ \boldsymbol{\beta}^{t,k_1\cdots k_t} \in \mathbb{R}_+^R,\ \boldsymbol{\gamma}^{t,k_1\cdots k_t} \in \mathbb{R}^{N_\xi},\ t \in \mathcal{T},\ k_1,\ldots,k_t \in \mathcal{K} \\
& \boldsymbol{y}^{t,k} \in \mathcal{Y}_t,\ \boldsymbol{w}^{t,k} \in \mathcal{W}_t \quad \forall t \in \mathcal{T},\ k \in \mathcal{K} \\
& \mathbf{e}^\top \boldsymbol{\alpha} = 1 \\
& \boldsymbol{A}^\top \boldsymbol{\beta}^{1,k_1} = \displaystyle\sum_{k_2\in\mathcal{K}} \boldsymbol{w}^{1,k_1} \circ \boldsymbol{\gamma}^{2,k_1k_2} \quad \forall k_1 \in \mathcal{K} \\
& \boldsymbol{A}^\top \boldsymbol{\beta}^{t,k_1\cdots k_t} + \boldsymbol{w}^{t-1,k_{t-1}} \circ \boldsymbol{\gamma}^{t,k_1\cdots k_t} = \displaystyle\sum_{k_{t+1}\in\mathcal{K}} \boldsymbol{w}^{t,k_t} \circ \boldsymbol{\gamma}^{t+1,k_1\cdots k_{t+1}} \quad \forall t \in \mathcal{T}\setminus\{1,T\},\ k_1,\ldots,k_t \in \mathcal{K} \\
& \boldsymbol{A}^\top \boldsymbol{\beta}^{T,k_1\cdots k_T} + \boldsymbol{w}^{T-1,k_{T-1}} \circ \boldsymbol{\gamma}^{T,k_1\cdots k_T} = \boldsymbol{\alpha}_{k_1\cdots k_T} \displaystyle\sum_{t\in\mathcal{T}} \left(\boldsymbol{D}^t\, \boldsymbol{w}^{t,k_t} + \boldsymbol{Q}^t\, \boldsymbol{y}^{t,k_t}\right) \quad \forall k_1,\ldots k_T \\
& \boldsymbol{w}^{t,k_t} \geq \boldsymbol{w}^{t-1,k_{t-1}} \quad \forall t \in \mathcal{T},\ k_t \in \mathcal{K},\ k_{t-1} \in \mathcal{K} \\
& \displaystyle\sum_{t\in\mathcal{T}} \boldsymbol{V}^t \boldsymbol{w}^{t,k_t} + \boldsymbol{W}^t \boldsymbol{y}^{t,k_t} \leq \boldsymbol{h} \quad \forall k_1,\ldots,k_T \in \mathcal{K}.
\end{array}
$$

that can be readily linearized using standard techniques.

8. Speed-Up Strategies & Extensions

This section proposes several strategies for speeding-up the solution of the K-adaptability counterpart of problems with exogenous and/or endogenous uncertainty. Then, it proposes an *approximate* method for generalizing the ideas in the paper to problems with continuous recourse.

8.1. Symmetry Breaking Constraints

The K-adaptability problem $(\mathcal{P}_K)$ presents a large amount of symmetry since indices of the candidate policies can be permuted to yield another, distinct, feasible solution with identical cost. This symmetry yields to significant slow down of the branch-and-bound procedure, see e.g., Bertsimas and Weismantel (2005), in particular as K grows. Thus, we propose to eliminate the symmetry in the problem by introducing symmetry breaking constraints. Specifically, we constrain the candidate policies $\{\boldsymbol{y}^k\}_{k\in\mathcal{K}}$ to be lexicographically decreasing. For this purpose, we introduce auxiliary binary variables $\boldsymbol{z}^{k,k+1} \in \{0,1\}^{N_y}$ for all $k \in \mathcal{K}\setminus\{K\}$

such that $\boldsymbol{z}_i^{k,k+1} = 1$ if and only if policies $\boldsymbol{y}^k$ and $\boldsymbol{y}^{k+1}$ differ in their ith component. These variables can be defined by means of a moderate number of linear inequality constraints, as follows

$$\left.\begin{aligned} \boldsymbol{z}_i^{k,k+1} &\leq \boldsymbol{y}_i^k + \boldsymbol{y}_i^{k+1} \\ \boldsymbol{z}_i^{k,k+1} &\leq 2 - \boldsymbol{y}_i^k - \boldsymbol{y}_i^{k+1} \\ \boldsymbol{z}_i^{k,k+1} &\geq \boldsymbol{y}_i^k - \boldsymbol{y}_i^{k+1} \\ \boldsymbol{z}_i^{k,k+1} &\geq \boldsymbol{y}_i^{k+1} - \boldsymbol{y}_i^k \end{aligned}\right\} \quad \forall i \in \mathcal{I},\ k \in \mathcal{K} \backslash \{K\}. \tag{23}$$

The first set of constraints above ensures that if $\boldsymbol{y}_i^k = \boldsymbol{y}_i^{k+1}$, then $\boldsymbol{z}_i^{k,k+1} = 0$. Conversely, the second set of constraints guarantees that $\boldsymbol{z}_i^{k,k+1} = 1$ whenever $\boldsymbol{y}_i^k \neq \boldsymbol{y}_i^{k+1}$. Using the variables $\boldsymbol{z}_i^{k,k+1}$, the lexicographic ordering constraints can be written as

$$\boldsymbol{y}_i^k \geq \boldsymbol{y}_i^{k+1} - \sum_{i' < i} \boldsymbol{z}_{i'}^{k,k+1} \quad \forall i \in \mathcal{I},\ k \in \mathcal{K} \backslash \{K\}. \tag{24}$$

These stipulate that if $\boldsymbol{y}_{i'}^k = \boldsymbol{y}_{i'}^{k+1}$ for all $i' < i$, then $\boldsymbol{y}_i^k \geq \boldsymbol{y}_i^{k+1}$. Since the symmetry breaking constraints in (23) and (24) are deterministic, they can be added to the K-adaptability problem without affecting the solution procedure.

8.2. The Case of Continuous Recourse Decisions

Throughout Sections 4, 5, and 6, we assume that both the here-and-now and wait-and-see decisions are binary. While the assumption that the here-and-now decisions are binary is not too restrictive, the assumption that the wait-and-see decisions are binary may be violated in many practical settings. It this section, we propose an approximation scheme that enables us to generalize our approach to the case of continuous recourse decisions.

Consider the following variant of Problem ($\mathcal{P}$) where the wait-and-see decisions $\boldsymbol{y}$ are real-valued ($\mathcal{Y} \subseteq \mathbb{R}^{N_y}$) and its coefficients in the objective function are deterministic.

$$\begin{aligned} \min \;\; & \max_{\overline{\boldsymbol{\xi}} \in \Xi} \; \min_{\boldsymbol{y} \in \mathcal{Y}} \left\{ \max_{\boldsymbol{\xi} \in \Xi(\boldsymbol{w}, \overline{\boldsymbol{\xi}})} \; \boldsymbol{\xi}^\top \boldsymbol{C}\, \boldsymbol{x} + \boldsymbol{\xi}^\top \boldsymbol{D}\, \boldsymbol{w} + \boldsymbol{q}^\top \boldsymbol{y} \;:\; \boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y} \leq \boldsymbol{H}\boldsymbol{\xi} \quad \forall \boldsymbol{\xi} \in \Xi(\boldsymbol{w}, \overline{\boldsymbol{\xi}}) \right\} \\ \text{s.t.} \;\; & \boldsymbol{x} \in \mathcal{X},\ \boldsymbol{w} \in \mathcal{W}, \end{aligned} \tag{25}$$

where $\boldsymbol{q} \in \mathbb{R}^{N_y}$. In the spirit of the linear decision rule approximation approach proposed in the stochastic and robust optimization literature, see e.g., Ben-Tal et al. (2004), Kuhn et al. (2009), Bodur and Luedtke (2018), we propose to restrict the recourse decisions $\boldsymbol{y}$ to those that are expressible as

$$\boldsymbol{y}(\boldsymbol{\xi}) = \boldsymbol{Y}\boldsymbol{\xi},$$

for some matrix $\boldsymbol{Y} \in \{0,1\}^{N_y \times N_\xi}$. This model is very natural since it enables us for example to choose between a large number of modes of operation for the wait-and-see decisions.

Under this approximation, Problem (25) is equivalent to

$$\begin{array}{ll} \text{minimize} & \displaystyle \max_{\overline{\boldsymbol{\xi}} \in \Xi(\boldsymbol{w})} \; \min_{\boldsymbol{Y} \in \{0,1\}^{N_y \times N_\xi}} \left\{ \begin{array}{ll} \displaystyle \max_{\boldsymbol{\xi} \in \Xi(\boldsymbol{w}, \overline{\boldsymbol{\xi}})} & \boldsymbol{\xi}^\top \boldsymbol{C}\, \boldsymbol{x} + \boldsymbol{\xi}^\top \boldsymbol{D}\, \boldsymbol{w} + \boldsymbol{q}^\top \boldsymbol{Y} \boldsymbol{\xi} \\ \text{s.t.} & \boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{Y}\boldsymbol{\xi} \leq \boldsymbol{H}\boldsymbol{\xi} \quad \forall \boldsymbol{\xi} \in \Xi(\boldsymbol{w}, \overline{\boldsymbol{\xi}}) \\ & \boldsymbol{Y}\boldsymbol{\xi} \in \mathcal{Y} \quad \forall \boldsymbol{\xi} \in \Xi(\boldsymbol{w}, \overline{\boldsymbol{\xi}}) \end{array} \right\} \\ \text{subject to} & \boldsymbol{x} \in \mathcal{X},\; \boldsymbol{w} \in \mathcal{W}. \end{array}$$

This problem is precisely of the form ($\mathcal{P}$) with the matrix $\boldsymbol{W}$ being affected by uncertainty (left-handside uncertainty). From Remark 6, our K-adaptability approximation framework applies in this case too. It results in a number K of linear contingency plans or *operating regimes*.

9. Robust Active Preference Learning at LAHSA

In this section, we propose two formulations of a preference elicitation problem that explicitly capture the endogenous nature of the elicitation process. In the first model, we take the point of view of a risk-averse decision-maker that seeks to maximize the *worst-case utility* of the item recommended at the end of the elicitation process. In the second model, we instead take the point of view of a regret averse decision-maker that wishes to minimize the *worst-case regret* of the offered item, see Section 6.4.

9.1. Motivation & Problem Formulation

The motivation for our study is the problem of designing a policy for allocating scarce housing resources to those experiencing homelessness that meets the needs of policy-makers at LAHSA, the lead agency in charge of public housing allocation in L.A. County. Thus, we consider the problem faced by a recommendation system which seeks to offer a user (in this case a policy-maker) with unknown preferences their favorite item (policy) among a finite but potentially large collection. Before making a recommendation, the system has the opportunity to elicit the user's preferences by making a moderate number of queries. Each query corresponds to a choice of an item (in this case a policy). The user is asked to respond with a number between 0 (zero) and 1 (one) with 0 (resp. 1) corresponding the *least* (resp. *most*) anyone could like an item. We take, in turn, the point of view of a risk-averse and of a regret-averse recommendation system which only possesses limited, set-based, information on the user utility function. Next, we introduce the notation in our model.

The main building blocks of our framework are candidate items (e.g., policies with different characteristics/outcomes) that the recommendation system can use to make queries or to offer to the user as a recommendation. We let $\mathcal{F} \subseteq \mathbb{R}^J$ denote the *universe* of all possible items. Each item $\boldsymbol{\phi} \in \mathcal{F}$ is uniquely characterized by its J features (or attributes), $\boldsymbol{\phi}_j \in \mathcal{F}_j \subseteq \mathbb{R}$, $j = 1, \ldots, J$, where $\mathcal{F}_j$ denotes the support of feature j. Thus, we use the feature values of an item to define the item. If two items have the same features, they are considered identical. We assume that $|\mathcal{F}|$ is finite and index items in the item set by $i \in \mathcal{I} := \{1, \ldots, I\}$. We denote by $\boldsymbol{\phi}^i \in \mathbb{R}^J$ the ith item in the query set, $i \in \mathcal{I}$. Thus, $\mathcal{F} = \{\boldsymbol{\phi}^i \mid i \in \mathcal{I}\}$.

We propose to represent user preferences with an (unknown) utility function $u : \mathcal{F} \to \mathbb{R}$ and analyze the user's behavior indirectly with utility functions. The function u ranks each item in the universe $\mathcal{F}$ and enables us to quantify how much the user prefers one item over another. We thus treat utilities as cardinal (as opposed to only ordinal) since they provide information on the strength of the preferences rather than just on the rank ordering of each item. We assume that the utility function is linear being defined through $u(\boldsymbol{\phi}) := \boldsymbol{u}^\top \boldsymbol{\phi}$ for some (unknown) vector $\boldsymbol{u}$ supported in the set $\mathcal{U} \subseteq [-1, 1]^J$. Note that the assumption that the utility function is linear in $\boldsymbol{\phi}$ and $\boldsymbol{u}$ is standard in the literature. Moreover, the requirement that $\boldsymbol{u}$ be normalized and bounded is without loss of generality since utility coefficients can be all scaled by a constant without affecting the relative strength of preferences over items.

Before recommending an item from the set $\mathcal{F}$, the recommendation system has the opportunity to make a number Q of queries to the user. These queries may enable the system to gain information about $\boldsymbol{u}$, thus improving the quality of the recommendation. Each query takes the form of an item: specifically, if item 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 item i?" The true answer to this question is given by the (normalized) quantity

$$\boldsymbol{\zeta}_i = \quad (\boldsymbol{u}^\top \boldsymbol{\phi}_i + \max_{j \in \mathcal{I}} \|\boldsymbol{\phi}_j\|_1)/(2 \max_{j \in \mathcal{I}} \|\boldsymbol{\phi}_j\|_1).$$

Indeed, note that the maximum utility that any user with utility $\boldsymbol{u} \in [-1, 1]^J$ can enjoy from an item is given by $\max_{j \in \mathcal{I}} \max_{\boldsymbol{u} \in [-1,1]^J} \boldsymbol{u}^\top \boldsymbol{\phi}_j = \max_{j \in \mathcal{I}} \|\boldsymbol{\phi}_j\|_1$. Similarly, the minimum utility that any user with utility $\boldsymbol{u} \in [-1, 1]^J$ can enjoy from an item is given by $\min_{j \in \mathcal{I}} \min_{\boldsymbol{u} \in [-1,1]^J} \boldsymbol{u}^\top \boldsymbol{\phi}_j = \min_{j \in \mathcal{I}} -\|\boldsymbol{\phi}_j\|_1$. Thus, the normalization proposed ensures that, for any feasible user and any feasible item, the quantity $\boldsymbol{\zeta}_i$ lies in the range $[0, 1]$. It also ensures that there is at least one utility vector resulting in an item with unit utility and, accordingly, that there is at least one utility vector resulting in an item with 0 utility.

Several authors have shown that oftentimes, individuals behave in seemingly "irrational" ways, see e.g., Kahneman and Tversky (1979), Kahneman et al. (1991), Allais (1953). In particular, they have shown that, when describing their preferences, users may give answers that are inconsistent and could be influenced by the framing of the question. In our framework, we cater for such inconsistencies explicitly. In particular, we assume that $\boldsymbol{\zeta}_i$ is not directly observable and that, even if the user is asked question i, we will only observe $\boldsymbol{\zeta}_i + \boldsymbol{\epsilon}_i$ where $\boldsymbol{\epsilon}_i$ is additive noise perturbing the answer to the question, i.e., we only observe a noisy version of the true normalized user utility for item i. We assume that the $\boldsymbol{\epsilon}_i$, $i \in \mathcal{I}$, are independent, identically distributed random variables with zero mean and given standard deviation. Then, in the spirit of modern robust optimization, see e.g., Lorca and Sun (2016), we assume that $\boldsymbol{\epsilon}$ is valued in the set $\mathcal{E} := \left\{ \boldsymbol{\epsilon} \in \mathbb{R}^I \,:\, \sum_{i=1}^I |\boldsymbol{\epsilon}_i| \leq \Gamma \right\}$, where Γ is a user-specified *budget of uncertainty* parameter that controls the degree of conservatism. The uncertainty set Ξ can be expressed as

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

Remark 7. *In the preference elicitation literature (see e.g., Boutilier et al. (2004), Bertsimas and O'Hair (2013)), robustness is usually achieved by asking the user* pairwise comparisons *(e.g., "do you prefer item A, or item B"). In the present paper, we propose an alternative way to achieve robustness (through the parameter* $\boldsymbol{\epsilon}$*).*

For each $i \in \mathcal{I}$, let $\boldsymbol{w}_i \in \{0,1\}$ denote the decision indicating whether $\boldsymbol{\xi}_i$ is observed today. Thus, $\boldsymbol{w}_i = 1$ if and only if we ask the user how much he likes item i. We assume that Q questions may be asked, i.e., $\mathcal{W} := \left\{ \boldsymbol{w} \in \{0,1\}^I \,:\, \sum_{i \in \mathcal{I}} \boldsymbol{w}_i = Q \right\}$. We let $\boldsymbol{y}_i \in \{0,1\}$ denote the decision to recommend item i, $i \in \mathcal{I}$, after the subset of elements of $\boldsymbol{\xi}$ that we chose to observe is revealed. We assume that only a single item can be recommended, i.e., $\mathcal{Y} := \left\{ \boldsymbol{y} \in \{0,1\}^I \,:\, \sum_{i \in \mathcal{I}} \boldsymbol{y}_i = 1 \right\}$.

With the above assumptions, we consider two variants of the active preference elicitation problem. In the first, we seek to select questions that will yield recommendations that have the highest worst-case (max-min) utility, solving the robust preference elicitation problem

$$\begin{array}{ll} \text{maximize} & \displaystyle \min_{\overline{\boldsymbol{\xi}} \in \Xi} \max_{\boldsymbol{y} \in \mathcal{Y}} \left\{ \min_{\boldsymbol{\xi} \in \Xi(\boldsymbol{w}, \overline{\boldsymbol{\xi}})} \boldsymbol{\xi}^\top \boldsymbol{y} \right\} \\ \text{subject to} & \boldsymbol{w} \in \mathcal{W}. \end{array} \qquad (\mathcal{WCU}^{\text{PE}})$$

Problem ($\mathcal{WCU}^{\text{PE}}$) is a two-stage robust problem with DDID and *objective uncertainty.* In the second variant, we seek to minimize the worst-case (min-max) absolute regret of the recommendation given by

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

see Section 6.4 and Observation 3. 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 $\boldsymbol{w}$. Problems ($\mathcal{WCU}^{\text{PE}}$) and ($\mathcal{WCAR}^{\text{PE}}$) can be solved approximately using the K-adaptability approximation schemes discussed in Sections 4 and 6, respectively. Indeed, the regret in Problem ($\mathcal{WCAR}^{\text{PE}}$) is given as the maximum of finitely many linear functions and Theorem 6 applies. We note that $|\mathcal{Y}| = I$ for both the worst-case utility and worst-case regret problems, ($\mathcal{WCU}^{\text{PE}}$) and ($\mathcal{WCAR}^{\text{PE}}$). Thus, solving the K-adaptability counterparts of ($\mathcal{WCU}^{\text{PE}}$) and ($\mathcal{WCAR}^{\text{PE}}$) with $K = I$ recovers an optimal solution to the corresponding original problem. Unfortunately, as I grows, solving the I-adaptability problem becomes prohibitive and, in practice, we are forced to focus on values of $K < I$, see Section 10. Yet, as we will see in our computational results in Section 10, highly competitive solutions can be identified even with such values of K.

10. Computational Studies

In this section, we investigate the performance of our approach on a variety of synthetic and real-world instances of the max-min utility and min-max regret preference elicitation problems introduced in Section 9. This section is organized at follows. In Section 10.1, we describe the datasets that we employ in our experiments. The experimental setup is described in Section 10.2. In Sections 10.3 and 10.4, we present our numerical results on the synthetic and real datasets, respectively.

10.1. Description of the Datasets

Synthetic Datasets. We generate a variety of synthetic datasets with $(I, J) \in \{10, 15, 20, 30, 40, 60\} \times \{10, 20, 30\}$. Each feature j of item i, $\boldsymbol{\phi}^i_j$, is drawn uniformly at random from the $[-1, 1]$ interval.

Real Dataset from LAHSA. The starting point of our analysis is the HMIS (Homeless Management Information System) dataset that our partners at LAHSA have made available to us and that pertains to the entire L.A. County. This dataset enables us to track the full trajectories of individuals in the public housing system. Specifically, for each individual waitlisted, it shows all the *supportive* housing resources (e.g., Permanent Supportive Housing –PSH–, Rapid Rehousing –RRH–) they received and the duration of their stay in each resource. It indicates temporary housing solutions they obtained (e.g., "Shelters") and any instances of individuals that exited homelessness without support (e.g., they "self-resolved" or returned permanently to their family). For each individual, it also shows several personal characteristics, e.g., gender, age, race, and answers to the Vulnerability Index-Service Prioritization Decision Assistance Tool[5] (VI-SPDAT) questions. Using this data, we use random forests to learn the probability that any given individual will exit homelessness if given a particular resource or on their own (service only). Then, we design 20 candidate parametric policies (in the form of linear or decision-tree based policies) for prioritizing individuals for housing. We proceed as follows for each candidate policy: *a)* we sample a type of policy from either logistic or decision-tree based; *b)* we sample a number of features from the data between 1 and 8; *c)* we learn, for each of PSH and RRH resources, and for service only, the probability that any given individual who gets the resource will successfully exit homelessness. Finally, we simulate the performance of the housing allocation system under each of these policies. For each policy, we record the following characteristics: *a)* the number of features it uses, and *b)* whether it is linear or decision-tree based, both of which serve as measures of interpretability; *c)* the probability that any given individual exits homelessness, overall, by race, and by gender; *d)* the average wait time, overall and by race; and *e)* whether the policy used protected features (e.g., race, gender, age), giving us a total of $J = 23$ features. Finally, we normalize the dataset by dividing each column by its infinity norm.

10.2. Experimental Setup

Throughout our experiments, and unless explicitly stated otherwise, the K-adaptability counterparts of Problems ($\mathcal{WCU}^{\rm PE}$) and ($\mathcal{WCAR}^{\rm PE}$) are solved using the techniques described in Sections 4 and 6, respectively. The tolerance δ used in the column-and-constraint generation algorithm (see Section 6.3) is 1×10^{-3}. We evaluate the *true* worst-case utility of any given solution $\boldsymbol{w}^\star$, which we denote by $u_{\rm wc}(\boldsymbol{w}^\star)$ as follows: we fix $\boldsymbol{w} = \boldsymbol{w}^\star$ in Problem (11) with $K = I$ (we explicitely set the I candidate policies to be all elements of $\mathcal{Y}$). Similarly, we evaluate the *true* worst-case regret of any given solution $\boldsymbol{w}^\star$, which we denote by $r_{\rm wc}(\boldsymbol{w}^\star)$, as follows: we fix $\boldsymbol{w} = \boldsymbol{w}^\star$ in Problem ($\mathcal{CCG}_{\rm feas}(\boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}^k\}_{k\in\mathcal{K}})$), where we set $K = I$ and employ all I candidate policies $\{\boldsymbol{y}^k\}_{k\in\mathcal{K}}$ in the set $\mathcal{Y}$.

To speed-up computation further, we leverage the structure of the preference elicitation problem to represent the symmetry breaking constraints (see Section 8.1) more efficiently, as discussed in Section EC.1.1. We also employ a conservative greedy heuristic that uses the solution to problems with smaller K to solve problems with larger K more efficiently, see Section EC.1.2. These strategies enable us to solve problems with large approximation parameters K (in the order of $K = 10$), which would not be solvable otherwise. We investigate the role played by these speed-up strategies at the end of Section 10.3. All of our experiments were performed on the High Performance Computing Cluster of our university. Each job was allotted 64GB of RAM, 16 cores, and a 2.6GHz Xeon processor. All optimization problems were solved using Gurobi version 8.0.0.

10.3. Numerical Results on Synthetic Data

Optimality-Scalability Trade-Off. In our first set of experiments, we evaluate the trade-off between computational complexity and scalability of our approach, as controlled by the single design parameter K. We also investigate the performance of our approach in terms of both optimality and scalability relative to the "prepartitioning" approach from Vayanos et al. (2011), see Section 3.3. As discussed in the introduction, this is the only approach that we are aware of that is directly applicable to two-stage robust problems with DDID. Our results on the synthetic datasets are summarized in Figures 3 and 4 for the max-min utility problem ($\mathcal{WCU}^{\rm PE}$) and the min-max regret problem ($\mathcal{WCAR}^{\rm PE}$), respectively. First, from the figures, we see that our proposed approach significantly outperforms the prepartitioning approach: indeed, as indicated by the position of the green star relative to the efficient frontier of the K-adaptability approach, our framework consistently dominated the prepartitioning approach in terms of both solver time and optimality across all our experiments. Second, we observe that the marginal benefit of increasing K decreases as K grows. Third, the benefit of employing the K-adaptability approach is more pronounced as the number of questions grows and as Γ grows. This is due to the need for more "adaptability" as the dimension of the uncertainty set grows. Finally, from the figures, it is apparent that the proposed K-adaptability approximation is more attractive from a usability perspective that the prepartitioning approach since it only requires tuning a single desing parameter (K) to trade-off between optimality and scalability.

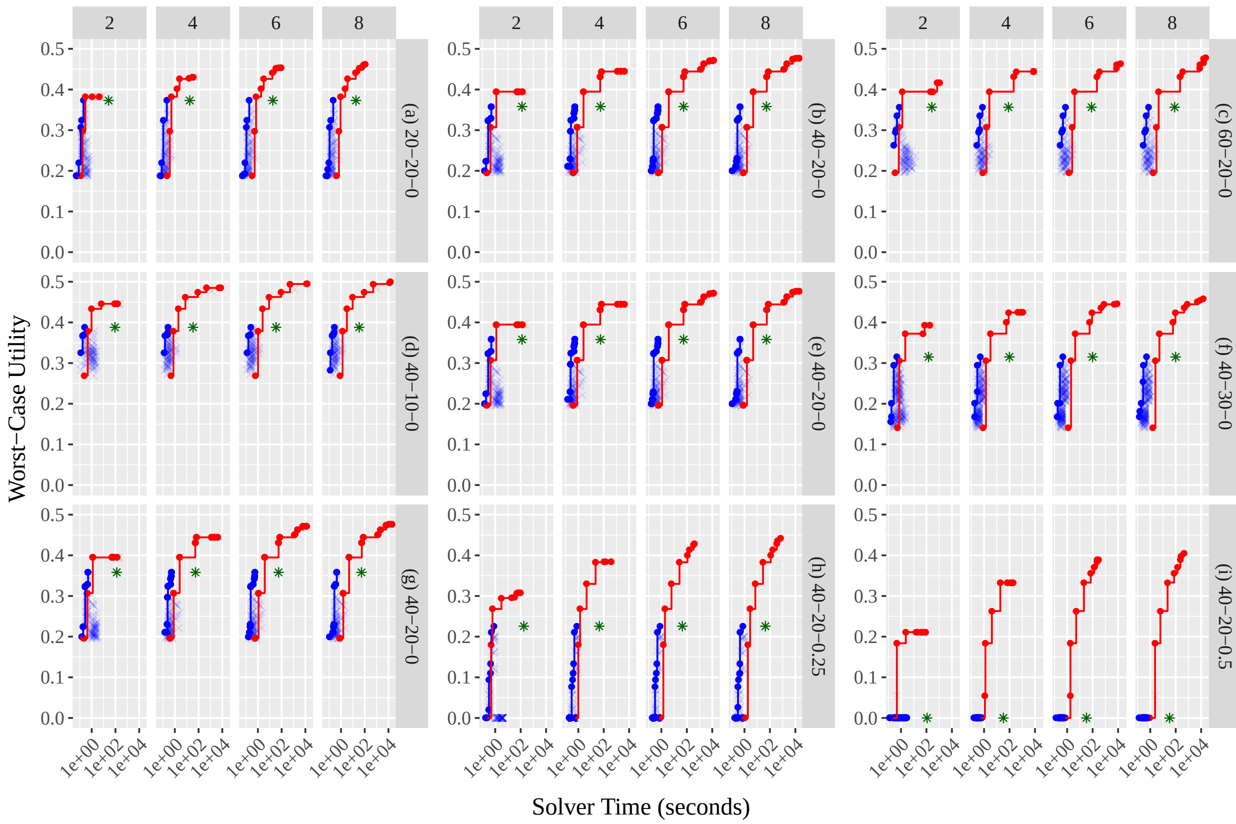

Figure 3 Optimality-scalability results for the max-min utility preference elicitation problem ($\mathcal{WCU}^{\text{PE}}$) on synthetic data. Each number on the top of each facet corresponds to the number of questions Q asked. Each label on the right of each facet corresponds to the characteristics of the instance solved ($I - J - \Gamma$). For example, the first four facets (subfigure (a)) are labeled 20–20–0, indicating an instance with 20 items, 20 features, and $\Gamma = 0$. On the first row (subfigures (a),(b), and (c)), we vary the number of items. On the second row (subfigures (d),(e), and (f)), we vary the number of features. On the last row (subfigures (g),(h), and (i)), we vary Γ. Each red dot (and cross) corresponds to a different choice of $K \in \{1, \ldots, 10\}$ for the K-adaptability problem. Each blue dot (and cross) corresponds to a different breakpoint configuration for the prepartitioning approach (we consider 100 different breakpoint configurations drawn randomly from the set of all configurations with cardinality less than 10). Whether a point is indicated with a dot or a cross depends on whether it is on the efficient frontier of the problems that resulted in the highest worst-case utility for the given time budget. The green star summarizes the performance of the prepartitioning approach: its solver time is calculated as the cumulative time needed to solve all prepartitioning problems and its worst-case utility corresponds to the best worst-case utility achievable by any breakpoint configuration.

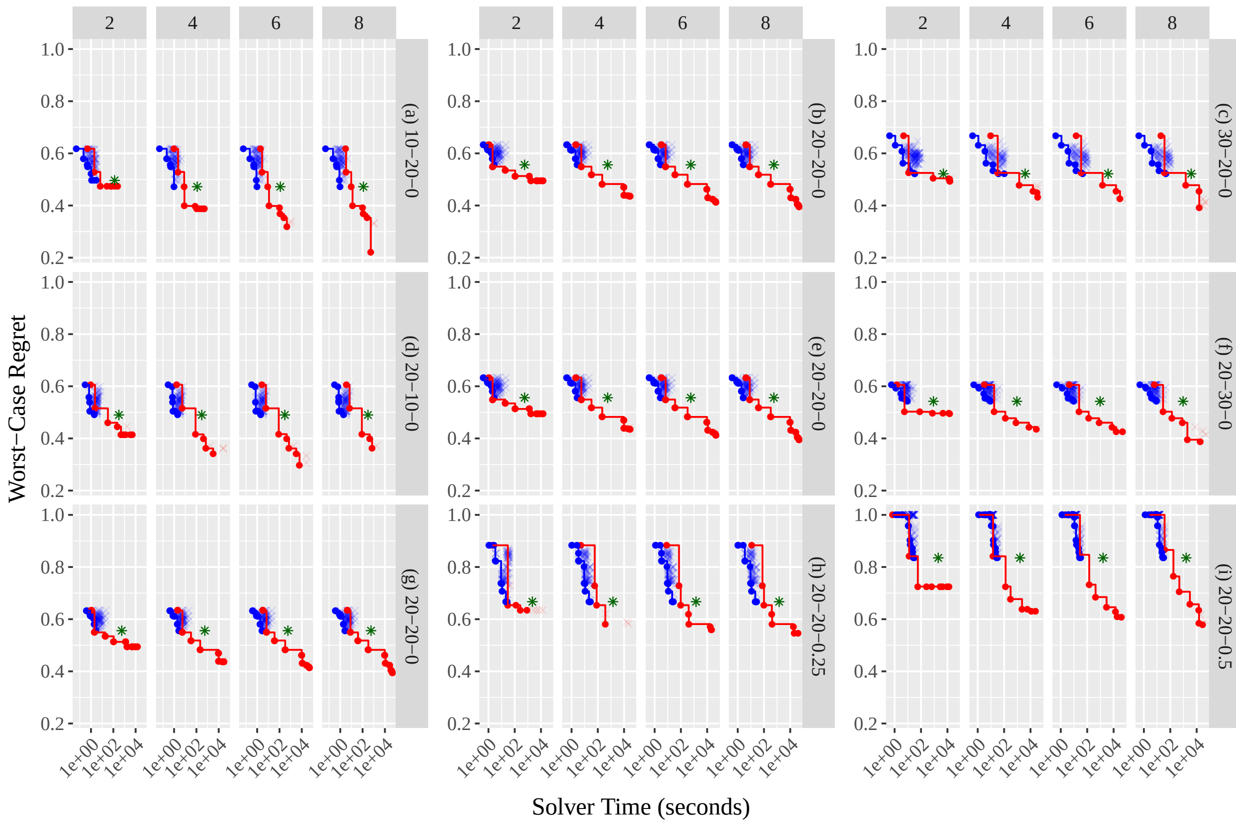

Figure 4 Optimality-scalability results for the min-max regret preference elicitation problem ($\mathcal{WCAR}^{\mathrm{PE}}$) on synthetic data. The facet labels, graphs, shapes, lines, and colors have the same interpretation as in Figure 3.

Performance Relative to Random Elicitation. In our second set of experiments, we evaluate the benefits of computing (near-)optimal queries using the K-adaptability approximation approach relative to asking questions at random. Thus, we compare the *true* performance of a solution to the K-adaptability problem to that of 1000 questions drawn uniformly at random from the set $\mathcal{W}$. Performance is measured in terms of the *true* objective value (true worst-case utility and true worst-case regret, respectively) of a given solution. The results are summarized on Figures 5 and 6 for the max-min utility and min-max regret cases, respectively. From the figures, it can be seen that the probability that the K-adaptability solution outperforms random elicitation converges to 0 as K grows. From Figure 5, we observe that in the case of the max-min utility problem, for values of K greater than approximately 5, the K-adaptability solution outperforms random elicitation in over 90% of the cases. That number increases to over 99% for K greater than about 8. From Figure 6, we observe that in the case of the min-max regret problem, higher values of about $K = 8$ are

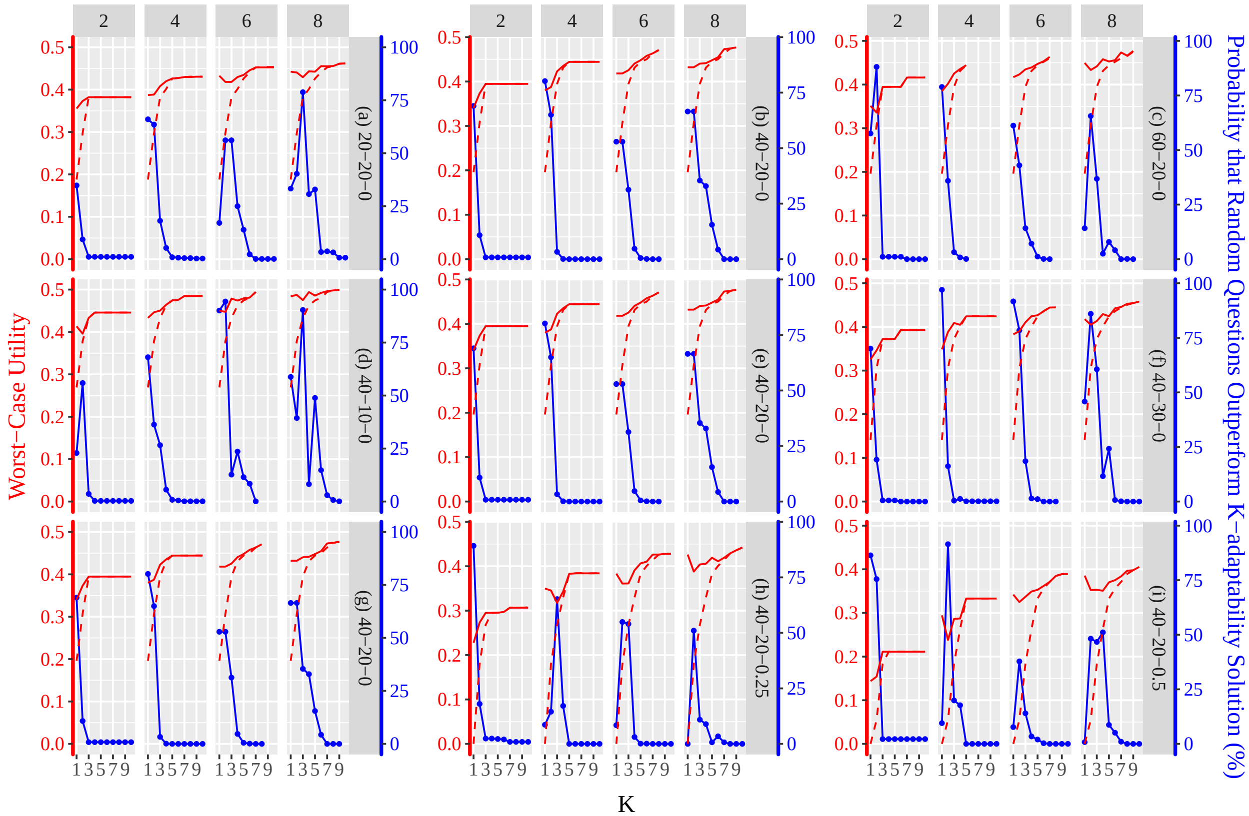

Figure 5 Results on the performance of the K-adaptability approach relative to random elicitation for the max-min utility preference elicitation problem ($\mathcal{WCU}^{\rm PE}$) on synthetic data. The dashed red line corresponds to the objective value of the K-adaptability problem. The red line corresponds to the true worst-case utility of the K-adaptability solution. The blue line represents the percentage of time that the true worst-case utility of a random solution exceeded the true worst-case utility of the K-adaptability solution.

needed to ensure that the K-adaptability solution outperforms random elicitation over 75% of the time. It is important to emphasize here that evaluating the true performance of a given solution is only possible for moderate problem sizes since it requires evaluating the objective value of the K adaptability counterpart for $K = I$, see Section 10.2. We also note that, to the best of our knowledge, evaluating $u_{\rm wc}(\boldsymbol{w}^\star)$ and $r_{\rm wc}(\boldsymbol{w}^\star)$ is only possible thanks to our proposed approach in this paper.

Comparison Between Max-Min Utility & Min-Max Regret Solutions. In the first two sets of experiments, we observed that the max-min utility problem, Problem ($\mathcal{WCU}^{\rm PE}$), is more scalable than its min-max regret counterpart, Problem ($\mathcal{WCAR}^{\rm PE}$). In our third set of experiments, we investigate whether there are benefits in employing the min-max regret solution relative the max-min utility solution. For this reason, we study

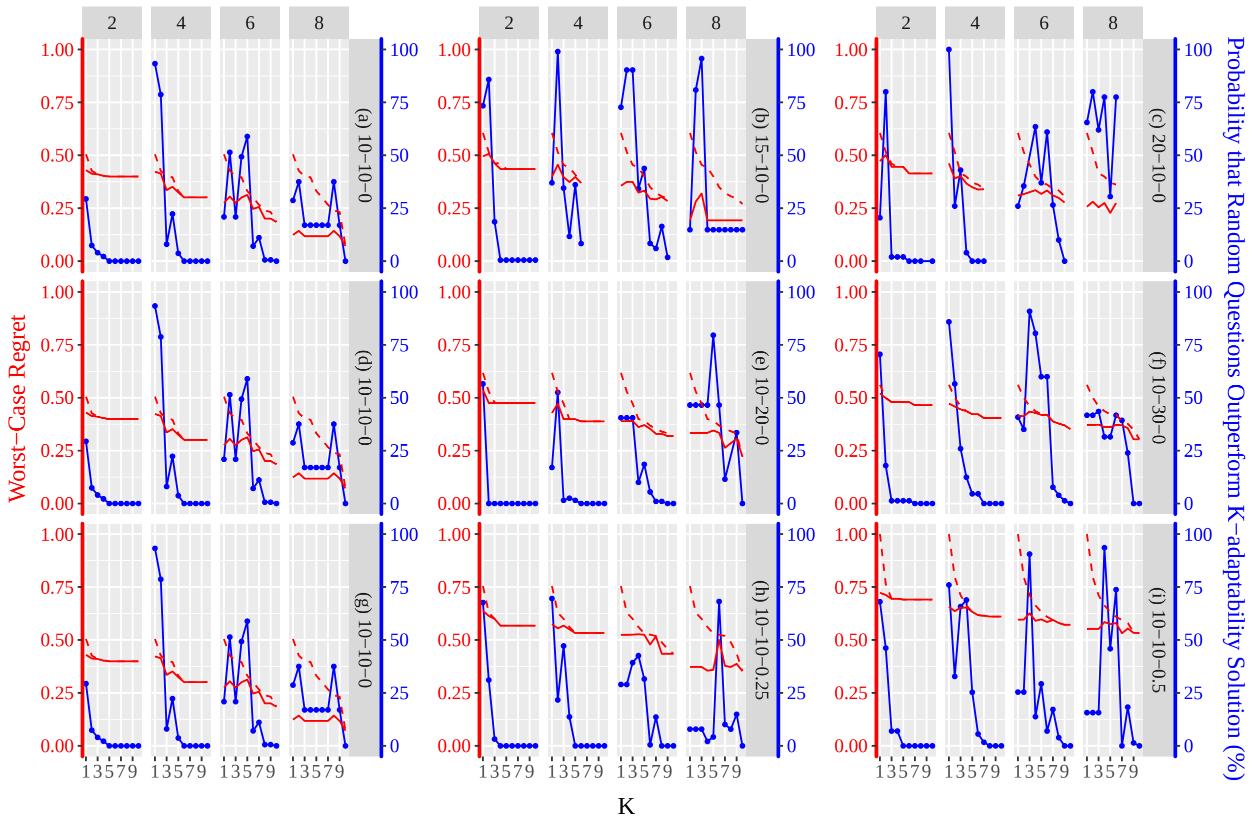

Figure 6 Results on the performance of the K-adaptability approach relative to random elicitation for the min-max regret preference elicitation problem ($\mathcal{WCAR}^{\mathrm{PE}}$) on synthetic data. The dashed red line corresponds to the objective value of the K-adaptability problem. The red line corresponds to the true worst-case regret of the K-adaptability solution. The blue line represents the percentage of time that the true worst-case regret of a random solution was lower that the true worst-case regret of the K-adaptability solution.

the true worst-case utility and true worst-case regret of solutions to the K-adaptability counterparts of Problems ($\mathcal{WCU}^{\mathrm{PE}}$) and ($\mathcal{WCAR}^{\mathrm{PE}}$), respectively, on a synthetic dataset with $I = 10$ items and $J = 10$ features ($\Gamma = 0$). The results are summarized in Table 1. From the table, it can be seen that, across all question budgets, employing the min-max regret criterion results in solutions that have far lower worst-case regret while simultaneously being competitive in terms of worst-case utility.

Evaluation of Symmetry Breaking & Greedy Heuristic. For our fourth set of experiments, we solved the max-min utility problem ($\mathcal{WCU}^{\mathrm{PE}}$) on a synthetic dataset with $I = 40$ items and $J = 20$ features ($\Gamma = 0$) using three different approaches: the K-adaptability counterpart, the K-adaptability counterpart augmented with symmetry breaking constraints, and the greedy heuristic approach, see Section EC.1.2. We varied K

Table 1 Comparison between max-min utility and min-max regret solutions on a synthetic dataset with $I=10$ items and $J=10$ features ($\Gamma=0$). The max-min utility solution $\boldsymbol{w}_{\mathrm{u}}^{\star}$ and min-max regret solution $\boldsymbol{w}_{\mathrm{r}}^{\star}$ are computed using the 10-adaptability counterparts of Problems ($\mathcal{WCU}^{\mathrm{PE}}$) and ($\mathcal{WCAR}^{\mathrm{PE}}$), respectively. The relative loss in true utility refers to the drop in worst-case utility experienced by employing the min-max regret rather than max-min utility solution, computed as $(u_{\mathrm{wc}}(\boldsymbol{w}_{\mathrm{u}}^{\star})-u_{\mathrm{wc}}(\boldsymbol{w}_{\mathrm{r}}^{\star}))/u_{\mathrm{wc}}(\boldsymbol{w}_{\mathrm{u}}^{\star})$. Similarly, the relative improvement in true regret refers to the improvement in worst-case regret experienced by employing the min-max regret rather than max-min utility solution, computed as $(r_{\mathrm{wc}}(\boldsymbol{w}_{\mathrm{r}}^{\star})-r_{\mathrm{wc}}(\boldsymbol{w}_{\mathrm{u}}^{\star}))/r_{\mathrm{wc}}(\boldsymbol{w}_{\mathrm{u}}^{\star})$.

Q	Max-Min Utility Solution		Min-Max Regret Solution		Relative Loss	Relative Improvement
	True Utility	True Regret	True Utility	True Regret	in True Utility	in True Regret
2	0.41	0.47	0.36	0.40	12.7%	14.8%
4	0.44	0.40	0.40	0.30	7.9%	24.3%
6	0.44	0.26	0.44	0.18	0.0%	29.9%
8	0.44	0.30	0.44	0.07	0.0%	76.5%

Table 2 Summary of evaluation results of symmetry breaking constraints and greedy heuristic approach on a synthetic dataset with $I=40$ items and $J=20$ features ($\Gamma=0$). The average speed-up factor in the solution of the max-min utility problem ($\mathcal{WCU}^{\mathrm{PE}}$) due to symmetry breaking constraints is computed by averaging over K the ratio of the solver time of the MILP without and with symmetry breaking constraints. Similarly, the average optimality gap of the greedy solution is computed by averaging over K the optimality gap of the greedy heuristic solution relative to the objective value of Problem ($\mathcal{WCU}^{\mathrm{PE}}$).

Number of Questions Q	2	4	6	8
Average Speed-Up Factor of Symmetry Breaking	13.2	4.7	4.3	7.4
Average Optimality Gap of Heuristic	5.4%	3.8%	3.8%	3.7%

from 1 to 10, and varied Q in the set $\{2,4,6,8\}$. The results are summarized on Table 2. From the table, we observe that the symmetry breaking constraints speed-up computation by a factor of 4.3 to 13.2 on average, depending on the number of questions asked. As seen in the table, the heuristic approach is near-optimal in all instances with a gap smaller than 6% in all cases. Detailed numerical results are provided in Section EC.7.

10.4. Numerical Results on the Real Dataset from LAHSA

Optimality-Scalability Trade-Off. In our fifth set of experiments, we evaluate the trade-off between computational complexity and scalability of our approach on the real dataset from LAHSA, see Section 10.1 for

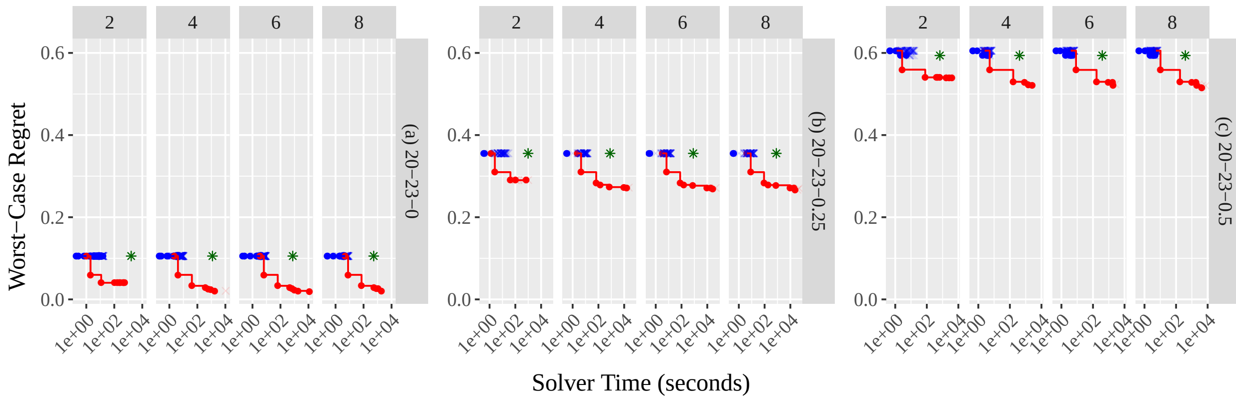

Figure 7 Optimality-scalability results for the min max regret preference elicitation problem ($\mathcal{WCAR}^{\rm PE}$) on the LAHSA data. The facet labels, graphs, shapes, lines, and colors have the same interpretation as in Figure 4.

details on this data. We solve both the max-min utility and min-max regret problems as Q and Γ are varied in the sets $\{2,4,6,8\}$ and $\{0,0.25,0.5\}$, respectively. The results are summarized in Figure 7. In this instance of the max-min utility problem, Problem ($\mathcal{WCU}^{\rm PE}$), static policies ($K=1$) are always optimal, with objective value 0.0265703, 0.0, and 0.0, for $\Gamma=0,0.25$, and 0.5, respectively. Thus, there is no benefit in asking any question and the performance of the K-adaptability and prepartitioning approaches are comparable. Intuitively, this is due to the fact that, in this dataset, the worst-case utility vector $\boldsymbol{u}$ remains unchanged after asking any one question. On the other hand, the K-adaptability approach significantly outperforms the prepartitioning approach in this instance of the min-max regret problem, Problem ($\mathcal{WCAR}^{\rm PE}$), and static policies are very sub-optimal (worst-case regret equal to 10.5%, 36%, and 60% for $\Gamma=0,0.25$, and 0.5, respectively). The prepartitioning approach performs comparably to static policy and only shows a small improvement in the case $\Gamma=0.5$. On the other hand, with the K-adaptability approach, the worst-case regret drops to 2.1%, 26.6%, and 51.4%, for $\Gamma=0,0.25$, and 0.5, respectively (for $Q=8$). Note that all the K-adaptability solutions to the min-max regret problem are optimal in the worst-case utility problem so the improvement in regret comes at no cost to worst-case utility. This experiment shows the strength of the K-adaptability approach; it also showcases the power of the min-max regret problem ($\mathcal{WCAR}^{\rm PE}$).

Performance Relative to Random Elicitation. In our final set of experiments, we evaluate the benefits of computing (near-)optimal queries using the K-adaptability approximation approach relative to asking questions at random on the real dataset from LAHSA. Similarly to the synthetic case, we compare the

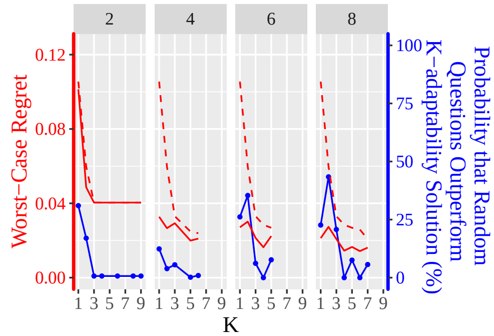

Figure 8 Results on the performance of the K-adaptability approach relative to random elicitation for the min-max regret preference elicitation problem ($\mathcal{WCAR}^{\mathrm{PE}}$) on the LAHSA dataset. The facet labels, graphs, shapes, lines, and colors have the same interpretation as in Figure 6.

true performance of a solution to the K-adaptability problem to that of 50 questions drawn uniformly at random from the set $\mathcal{W}$. We only consider a moderate number of samples in this case since evaluating the true regret on this larger dataset is computationally intractable. We measure performance in terms of the *true* worst-case regret of a given solution. The results are summarized on Figure 8. From the figure, it can be seen that the probability that the K-adaptability solution outperforms random elicitation converges to 0 as K grows. We observe that, for values of K greater than approximately 4, the K-adaptability solution outperforms random elicitation in over 90% of the cases.

Notes

[1]See `https://www.mturk.com/`.

[2]See `https://www.hudexchange.info/programs/hmis/`.

[3]See `https://www.lahsa.org/`.

[4]See e.g., `https://www.ibm.com/analytics/cplex-optimizer` and `https://www.gurobi.com/`.

[5]See `https://www.orgcode.com/`.

References

Adida E, Perakis G (2006) A robust optimization approach to dynamic pricing and inventory control with no backorders. *Mathematic Programming* 107(1):97–129.

Allais M (1953) Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'ecole americaine. *Econometrica* 21(4):503.

Ardestani-Jaafari A, Delage E (2016) Robust optimization of sums of piecewise linear functions with application to inventory problems. *Operations Research* 64(2):474–494.

Artstein Z, Wets RJB (1993) Sensors and information in optimization under stochastic uncertainty. Technical Report 3.

Assavapokee T, Realff MJ, Ammons JC (2008a) Min-max regret robust optimization approach on interval data uncertainty. *Journal of Optimization Theory and Applications* .

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. *Computers and Operations Research* .

Atamtürk A, Zhang M (2007) Two-stage robust network flow and design under demand uncertainty. *Operations Research* 55(4):662–673.

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

Bandi C, Trichakis N, Vayanos P (2018) Robust multiclass queuing theory for wait time estimation in resource allocation systems. *Management Science* 65(1):152–187.

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

Ben-Tal A, Golany B, Nemirovski A, Vial JP (2005) Retailer-supplier flexible commitments contracts: a robust optimization approach. *Manufacturing & Service Operations Management* 7(3):248–271.

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

Ben-Tal A, Nemirovski A (1998) Robust convex optimization. *Mathematics of Operations Research* 23(4):769–805.

Ben-Tal A, Nemirovski A (1999) Robust solutions of uncertain linear programs. *Operations Research Letters* 25:1–13.

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

Bertsimas D, Brown D, Caramanis C (2010) Theory and applications of robust optimization. *SIAM Review* 53(3):464–501.

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

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

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

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

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

Bertsimas D, Goyal V (2012) On the power and limitations of affine policies in two-stage adaptive optimization problems. *Mathematical Programming, Series A* .

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

Bertsimas D, O'Hair A (2013) Learning Preferences Under Noise and Loss Aversion: An Optimization Approach. *Operations Research* 61(5):1190–1199.

Bertsimas D, Pachamanova D, Sim M (2004) Robust linear optimization under general norms. *Operations Research Letters* 32(6):510–516.

Bertsimas D, Shtern S (2017) A scalable primal-dual algorithm for two-stage adaptive mixed-integer optimization. Technical report.

Bertsimas D, Sim M (2003) Robust discrete optimization and network flows. *Mathematical Programming* 98(1):49–71.

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

Bertsimas D, Vayanos P (2017) Data-driven learning in dynamic pricing using adaptive robust optimization.

Bertsimas D, Weismantel R (2005) *Optimization over integers.* (Athena Scientific), ISBN 978-0-97591-462-5.

Bodur M, Luedtke JR (2018) Two-stage linear decision rules for multi-stage stochastic programming. *Mathematical Programming* .

Boutilier C, Sandholm T, Shields R (2004) Eliciting Bid Taker Non-Price Preferences in (Combinatorial) Auctions. 204–211.

Chan TCY, Shen ZJM, Siddiq A (2018) Robust defibrillator deployment under cardiac arrest location uncertainty via row-and-column generation. *Operations Research* 66(2):358–379.

Chassein A, Goerigk M, Kurtz J, Poss M (2019) Faster algorithms for min-max-min robustness for combinatorial problems with budgeted uncertainty. *European Journal of Operational Research* 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 Transactions on Power Systems* .

Colvin M, Maravelias CT (2008) A stochastic programming approach for clinical trial planning in new drug development. *Computers & Chemical Engineering* 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. *European Journal of Operational Research* 203(1):205–215.

Gabrel V, Murat C, Thiele A (2014) Recent advances in robust optimization: An overview.

Georghiou A, Kuhn D, Wiesemann W (2018) The decision rule approach to optimization under uncertainty: methodology and applications. *Computational Management Science* .

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

Goel V, Grossman IE (2006) A class of stochastic programs with decision dependent uncertainty. *Mathematical 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. *Computers & Chemical Engineering* 30(6–7):1076–1092.

Goldfarb D, Iyengar G (2003) Robust portfolio selection problems. *Mathematics of Operations Research* 28:1–38.

Goldfarb D, Iyengar G (2004) Robust portfolio management. Technical report.

Gorissen B, Yanikoglu I, den Hertog D (2015) A practical guide to robust optimization. *Omega* 53:124–137.

Gounaris CE, Wiesemann W, Floudas CA (2013) The Robust Capacitated Vehicle Routing Problem Under Demand Uncertainty. *Operations Research* 61(3):677–693.

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

Gupta V, Han BR, Kim SH, Paek H (2017) Maximizing intervention effectiveness. *SSRN preprint* .

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

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

Hillier FS (2012) *Introduction To Operations Research* (McGraw-Hill Education (India) Pvt Limited), ISBN 9780071333467.

Jiang R, Wang J, Zhang M, Guan Y (2013) Two-stage minimax regret robust unit commitment. *IEEE Transactions on Power Systems* .

Jiang R, Zhang M, Li G, Guan Y (2014) Two-stage network constrained robust unit commitment problem. *European Journal of Operational Research* 234(3):751–762.

Jonsbråten TW (1998) *Optimization models for petroleum field exploitation*. Ph.D. thesis, Norwegian Shool of Economics and Business Administration.

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

Kahneman D, Knetsch JL, Thaler RH (1991) Anomalies: the endowment effect, loss aversion, and status quo bias. *Journal of Economic Perspectives* 5(1):193–206.

Kahneman D, Tversky A (1979) Prospect Theory: an analysis of decision under risk. *Econometrica* 47(2):263.

Kuhn D, Wiesemann W, Georghiou A (2009) Primal and dual linear decision rules in stochastic and robust optimization. *Mathematical Programming* 130(1):177–209.

Lappas NH, Gounaris CE (2016) Multi-stage adjustable robust optimization for process scheduling under uncertainty. *AIChE Journal* 62(5).

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

Mamani H, Nassiri S, Wagner MR (2017) Closed-form solutions for robust inventory management. *Management Science* 63(5):1625–1643.

McElfresh D, Vayanos P, Dickerson J, Rice E (2019) Robust active preference learning. Working Paper to be submitted to Management Science, 2019 (non-archival version under review at the WINE conference).

Ng TS (2013) Robust regret for uncertain linear programs with application to co-production models. *European Journal of Operational Research* 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. *Computers & Chemical Engineering* 108:425–447.

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

Nohadani O, Sharma K (2016) Optimization under decision-dependent uncertainty. Technical report, Northwestern.

Postek K, Den Hertog D (2016) Multistage adjustable robust mixed-integer optimization via iterative splitting of the uncertainty set. *INFORMS Journal on Computing* 28(3).

Poursoltani M, Delage E (2019) Adjustable robust optimization reformulations of two-stage worst-case regret minimization problems. Technical report.

Rahmattalabi A, Vayanos P, Fulginiti A, Rice E, Wilder B, Yadav A, Tambe M (2019) Exploring algorithmic fairness in robust graph covering problems. Forthcoming In Proceedings of the 33rd Conference on Neural Information Processing Systems (NeurIPS).

Savage LJ (1951) The theory of statistical decision. *Journal of the American Statistical association* 46(253):55–67.

Shapiro A (2017) Interchangeability principle and dynamic equations in risk averse stochastic programming. *Operations Research Letters* 45(4):377–381.

Solak S, Clarke JP, Johnson EL, Barnes ER (2010) Optimization of R&D project portfolios under endogenous uncertainty. *European Journal of Operational Research* 207(1):420–433.

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

Subramanyam A, Gounaris CE, Wiesemann W (2017) K-Adaptability in two-stage mixed-integer robust optimization. *arXiv preprint arXiv:1706.07097* .

Thiele A (2009) Multi-product pricing via robust optimization. *Journal of Revenue and Pricing Management* 8:67–80.

Vayanos P, Kuhn D, Rustem B (2011) Decision rules for information discovery in multi-stage stochastic programming. *Proceedings of the 50th IEEE Conference on Decision and Control*, 7368–7373.

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

Wilder B, Yadav A, Immorlica N, Rice E, Tambe M, Wilder B, Yadav A, Immorlica N, Rice E, Tambe M (2017) Uncharted but not uninfluenced: influence maximization with an uncertain network. *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems* .

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

Yanikoglu I, Gorissen B, den Hertog D (2017) Adjustable robust optimization – A survey and tutorial.

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

Zhang M (2011) Two-stage minimax regret robust uncapacitated lot-sizing problems with demand uncertainty. *Operations Research Letters* .

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

Zhao C, Wang J, Watson J, Guan Y (2013) Multi-stage robust unit commitment considering wind and demand response uncertainties. {*IEEE*} *Transactions on Power Systems* 28(3):2708–2717.

Zhen J, den Hertog D, Sim M (2016) Adjustable robust optimization via Fourier-Motzkin elimination.

E-Companion

EC.1. Strategies for Speeding-up Computations in Section 10

In this section, we detail several strategies for speeding-up computations that we have employed in our analysis in Section 10.

EC.1.1. More Efficient Representation of the Symmetry Breaking Constraints

In Section 8.1, we provided symmetry breaking constraints applicable to any K-adaptability problem. The active learning problems presented in Section 9 possess an attractive structure in that any feasible recourse decision (a recommended item) must have exactly one non-zero element. This observation enables us to enforce lexicographic ordering of the policies without needing auxiliary binary variables $\boldsymbol{z}^{k,k+1}$, which helps substantially reduce the size of the resulting problem and speed-up computation, see Section 10.

First, note that since having duplicate candidate policies does not improve the objective value and since $|\mathcal{Y}| = I$, duplicate policies can be excluded (thereby strengthening the associated formulation) by means of the constraints

$$\sum_{k\in\mathcal{K}} \boldsymbol{y}_i^k \leq 1 \quad \forall i \in \mathcal{I}, \tag{EC.1}$$

provided $K \leq I$. Second, provided constraints (EC.1) are imposed, the lexicographic ordering constraints reduce to

$$\boldsymbol{y}_i^k \geq \boldsymbol{y}_i^{k+1} - \sum_{i'<i} \boldsymbol{y}_{i'}^k - \sum_{i'<i} \boldsymbol{y}_{i'}^{k+1} \quad \forall i \in \mathcal{I},\, k \in \mathcal{K}\backslash\{K\}. \tag{EC.2}$$

This follows from the fact that, if (EC.1) holds, then policies $\boldsymbol{y}_{i'}^k$ and $\boldsymbol{y}_{i'}^{k+1}$ are identical for all $i' < i$ if and only if $\boldsymbol{y}_{i'}^k = \boldsymbol{y}_{i'}^{k+1} = 0$ for all $i' < i$. Indeed, if there exists $i' < i$ such that either $\boldsymbol{y}_{i'}^k = 1$ or $\boldsymbol{y}_{i'}^{k+1} = 1$, then it must hold that $\boldsymbol{y}_{i'}^k \neq \boldsymbol{y}_{i'}^{k+1}$. Thus, $\boldsymbol{y}_{i'}^k$ and $\boldsymbol{y}_{i'}^{k+1}$ are identical for all $i' < i$ if and only if $\sum_{i'<i} \boldsymbol{y}_{i'}^k + \sum_{i'<i} \boldsymbol{y}_{i'}^{k+1} = 0$, in which case we require $\boldsymbol{y}_i^k \geq \boldsymbol{y}_i^{k+1}$. This is precisely the constraint imposed in (EC.2). Adding the symmetry breaking constraints (EC.1) and (EC.2) to the K-adaptability counterpart of Problem ($\mathcal{WCU}^{\rm PE}$) or Problem ($\mathcal{WCAR}^{\rm PE}$) breaks the symmetry in the candidate policies. At the same time, it results in a far more efficient formulation (smaller number of decision variables and constraints) than adding the generic symmetry breaking constraints from Section 8.1.

Algorithm 2: Heuristic algorithm for solving the K-adaptability counterpart of a problem; adapted from Subramanyam et al. (2017).

Inputs: Instance of Problem ($\mathcal{PO}$), ($\mathcal{P}$), or ($\mathcal{PO}^{\rm PWL}$); K-adaptability parameter K;

Output: Conservative solution $(\boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}^k\}_{k\in\mathcal{K}})$ to the K-adaptability counterpart of the input instance (($\mathcal{PO}_K$), ($\mathcal{P}_K$), or ($\mathcal{PO}_K^{\rm PWL}$), respectively);

for $k \in \{1,\ldots,K\}$ **do**

  if $k=1$ **then**

    Solve the the k-adaptability counterpart of the input instance (using its MBLP reformulation);

    Let $(\boldsymbol{x}^\star, \boldsymbol{w}^\star, \boldsymbol{y}^{\star,1})$ denote an optimal solution;

  else

    Solve the the k-adaptability counterpart of the input instance (using its MBLP reformulation) with the added constraints that $\boldsymbol{y}^\kappa = \boldsymbol{y}^{\star,\kappa}$ for all $\kappa \in \{1,\ldots,k-1\}$;

    Let $(\boldsymbol{x}^\star, \boldsymbol{w}^\star, \{\boldsymbol{y}^{\star,\kappa}\}_{\kappa=1}^k)$ denote an optimal solution;

  end

end

Result: Return $(\boldsymbol{x}^\star, \boldsymbol{w}^\star, \{\boldsymbol{y}^{\star,\kappa}\}_{\kappa\in\mathcal{K}})$.

EC.1.2. Heuristic K-Adaptability Solution Approach

In addition to solving the K-adaptability counterpart of the active preference learning problems ($\mathcal{WCU}^{\rm PE}$) and ($\mathcal{WCAR}^{\rm PE}$) directly (using their MBLP reformulation), we also employ a heuristic approach, as detailed in Algorithm 2. A variant of this approach has been previously used by Subramanyam et al. (2017). This algorithm returns a feasible but potentially suboptimal solution to the K-adaptability counterpart of the problem to be solved.

EC.2. Proofs of Statements in Sections 2 and 3

Proof of Theorem 1 Let $\boldsymbol{x}$ and $\boldsymbol{y}(\cdot)$ be defined as in the premise of claim *(i)*. Then, $\boldsymbol{x} \in \mathcal{X}$ and, for each $\boldsymbol{\xi} \in \Xi$, we have that $\boldsymbol{y}(\boldsymbol{\xi}) \in \mathcal{Y}$ and $\boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}(\boldsymbol{\xi}) \leq \boldsymbol{H}\boldsymbol{\xi}$. Thus, $(\boldsymbol{x}, \boldsymbol{y}(\cdot))$ is feasible in Problem (1). Moreover, it is readily checked that the objective value $(\boldsymbol{x}, \boldsymbol{y}(\cdot))$ attains in Problem (1) is equal to the objective value attained by $\boldsymbol{x}$ in Problem (2). We have thus shown that Problem (1) lower bounds Problem (2) and that if $\boldsymbol{x}$

is optimal in Problem (2), then the pair $(\boldsymbol{x}, \boldsymbol{y}(\cdot))$ is feasible in Problem (1) with the two solutions attaining the same cost in their respective problems.

Next, let $(\boldsymbol{x}, \boldsymbol{y}(\cdot))$ be defined as in the premise of claim *(ii)*, i.e., let it be optimal in Problem (1). The here-and-now decision $\boldsymbol{x}$ is feasible in Problem (2) and, for each $\boldsymbol{\xi} \in \Xi$, we can define

$$\boldsymbol{y}'(\boldsymbol{\xi}) \in \operatorname*{arg\,min}_{\boldsymbol{y} \in \mathcal{Y}} \left\{ \boldsymbol{\xi}^\top \boldsymbol{C}\, \boldsymbol{x} + \boldsymbol{\xi}^\top \boldsymbol{Q}\, \boldsymbol{y} \,:\, \boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y} \leq \boldsymbol{H}\boldsymbol{\xi} \right\}.$$

By construction, $\boldsymbol{x} \in \mathcal{X}$. Moreover, by definition of $\boldsymbol{y}'(\cdot)$ and by feasibility of $(\boldsymbol{x}, \boldsymbol{y}(\cdot))$ in Problem (1) it holds that

$$\max_{\boldsymbol{\xi} \in \Xi} \;\; \boldsymbol{\xi}^\top \boldsymbol{C}\, \boldsymbol{x} + \boldsymbol{\xi}^\top \boldsymbol{Q}\, \boldsymbol{y}'(\boldsymbol{\xi}) \leq \max_{\boldsymbol{\xi} \in \Xi} \;\; \boldsymbol{\xi}^\top \boldsymbol{C}\, \boldsymbol{x} + \boldsymbol{\xi}^\top \boldsymbol{Q}\, \boldsymbol{y}(\boldsymbol{\xi}).$$

Thus, $\boldsymbol{x}$ is feasible in Problem (2) with a cost no greater than that of $(\boldsymbol{x}, \boldsymbol{y}(\cdot))$ in Problem (1). We have thus shown that Problem (2) lower bounds Problem (1) and that if $(\boldsymbol{x}, \boldsymbol{y}(\cdot))$ is optimal in Problem (1), then $\boldsymbol{x}$ is feasible in Problem (2) with the cost attained by $\boldsymbol{x}$ in Problem (2) being no greater than the cost of $(\boldsymbol{x}, \boldsymbol{y}(\cdot))$ in Problem (1).

We conclude that the optimal costs of Problems (1) and (2) are equal, and that claims *(i)* and *(ii)* hold. $\square$

Proof of Theorem 2 Let $(\boldsymbol{x}, \boldsymbol{w})$, $\boldsymbol{y}'(\cdot)$, and $\boldsymbol{y}(\cdot)$ be defined as in the premise of claim *(i)*. Then, $\boldsymbol{x} \in \mathcal{X}$, $\boldsymbol{w} \in \mathcal{W}$, and for each $\boldsymbol{\delta}$ such that $\boldsymbol{\delta} = \boldsymbol{w} \circ \overline{\boldsymbol{\xi}}$ for some $\overline{\boldsymbol{\xi}} \in \Xi$, we have that $\boldsymbol{y}'(\boldsymbol{\delta}) \in \mathcal{Y}$ and $\boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}'(\boldsymbol{\delta}) \leq \boldsymbol{H}\boldsymbol{\xi}$ for all $\boldsymbol{\xi} \in \Xi(\boldsymbol{w}, \boldsymbol{\delta})$. We show that $(\boldsymbol{x}, \boldsymbol{w}, \boldsymbol{y}(\cdot))$ is feasible in Problem (3). Fix any $\boldsymbol{\xi} \in \Xi$. First, $\boldsymbol{y}(\boldsymbol{\xi}) \in \mathcal{Y}$. Second, we have

$$\boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}(\boldsymbol{\xi}) \;=\; \boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}'(\boldsymbol{w} \circ \boldsymbol{\xi}) \;\leq\; \boldsymbol{H}\boldsymbol{\xi},$$

where the equality follows by definition of $\boldsymbol{y}(\cdot)$ and the inequality follows from the fact that $\boldsymbol{\xi} \in \Xi(\boldsymbol{w}, \boldsymbol{w} \circ \boldsymbol{\xi})$ and from the definition of $\boldsymbol{y}'(\cdot)$. Fix $\boldsymbol{\xi}' \in \Xi : \boldsymbol{w} \circ \boldsymbol{\xi} = \boldsymbol{w} \circ \boldsymbol{\xi}'$. Then, $\boldsymbol{y}(\boldsymbol{\xi}) = \boldsymbol{y}(\boldsymbol{\xi}')$, so that the decision-dependent non-anticipativity constraints are also satisfied. Since the choice of $\boldsymbol{\xi} \in \Xi$ was arbitrary, $(\boldsymbol{x}, \boldsymbol{w}, \boldsymbol{y}(\cdot))$ is feasible in Problem (3). The objective value attained by $(\boldsymbol{x}, \boldsymbol{w})$ in Problem $(\mathcal{P})$ is given by

$$\max_{\substack{\overline{\boldsymbol{\xi}} \in \Xi, \\ \boldsymbol{\xi} \in \Xi(\boldsymbol{w}, \overline{\boldsymbol{\xi}})}} \;\; \boldsymbol{\xi}^\top \boldsymbol{C}\, \boldsymbol{x} + \boldsymbol{\xi}^\top \boldsymbol{D}\, \boldsymbol{w} + \boldsymbol{\xi}^\top \boldsymbol{Q}\, \boldsymbol{y}'(\boldsymbol{w} \circ \overline{\boldsymbol{\xi}}) \;=\; \max_{\substack{\overline{\boldsymbol{\xi}} \in \Xi, \\ \boldsymbol{\xi} \in \Xi(\boldsymbol{w}, \overline{\boldsymbol{\xi}})}} \;\; \boldsymbol{\xi}^\top \boldsymbol{C}\, \boldsymbol{x} + \boldsymbol{\xi}^\top \boldsymbol{D}\, \boldsymbol{w} + \boldsymbol{\xi}^\top \boldsymbol{Q}\, \boldsymbol{y}(\boldsymbol{\xi}), \tag{EC.3}$$

where we have grouped the two maximization problems in a single one and where the equality follows from the definition of $\boldsymbol{y}(\cdot)$. The value attained by $(\boldsymbol{x}, \boldsymbol{w}, \boldsymbol{y}(\cdot))$ in Problem (3) is

$$\max_{\boldsymbol{\xi} \in \Xi} \;\; \boldsymbol{\xi}^\top \boldsymbol{C}\, \boldsymbol{x} + \boldsymbol{\xi}^\top \boldsymbol{D}\, \boldsymbol{w} + \boldsymbol{\xi}^\top \boldsymbol{Q}\, \boldsymbol{y}(\boldsymbol{\xi}). \tag{EC.4}$$

Since $\left\{\boldsymbol{\xi} \in \Xi(\boldsymbol{w},\overline{\boldsymbol{\xi}}) : \overline{\boldsymbol{\xi}} \in \Xi\right\} = \Xi$, it follows that the optimal objective values of the Problems (EC.3) and (EC.4) are equal. We have thus shown that Problem (3) lower bounds Problem ($\mathcal{P}$) and that if $(\boldsymbol{x},\boldsymbol{w})$ is optimal in Problem ($\mathcal{P}$), then the triple $(\boldsymbol{x},\boldsymbol{w},\boldsymbol{y}(\cdot))$ is feasible in Problem (3) with the two solutions attaining the same cost in their respective problems.

Next, let $(\boldsymbol{x},\boldsymbol{w},\boldsymbol{y}(\cdot))$ be defined as in the premise of claim *(ii)*, i.e., let it be optimal in Problem (3). The here-and-now decision $(\boldsymbol{x},\boldsymbol{w})$ is feasible in Problem ($\mathcal{P}$) and, for each $\overline{\boldsymbol{\xi}} \in \Xi$, we can define

$$\boldsymbol{y}'(\overline{\boldsymbol{\xi}}) \in \operatorname*{arg\,min}_{\boldsymbol{y}\in\mathcal{Y}} \left\{ \max_{\boldsymbol{\xi}\in\Xi(\boldsymbol{w},\overline{\boldsymbol{\xi}})} \; \boldsymbol{\xi}^\top \boldsymbol{C}\,\boldsymbol{x} + \boldsymbol{\xi}^\top \boldsymbol{D}\,\boldsymbol{w} + \boldsymbol{\xi}^\top \boldsymbol{Q}\,\boldsymbol{y} \;:\; \boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y} \le \boldsymbol{H}\boldsymbol{\xi} \quad \forall \boldsymbol{\xi} \in \Xi(\boldsymbol{w},\overline{\boldsymbol{\xi}}) \right\}.$$

By construction, $(\boldsymbol{x},\boldsymbol{w}) \in \mathcal{X} \times \mathcal{W}$. Moreover, it holds that

$$\begin{aligned} & \max_{\overline{\boldsymbol{\xi}}\in\Xi} \; \max_{\boldsymbol{\xi}\in\Xi(\boldsymbol{w},\overline{\boldsymbol{\xi}})} \; \boldsymbol{\xi}^\top \boldsymbol{C}\,\boldsymbol{x} + \boldsymbol{\xi}^\top \boldsymbol{D}\,\boldsymbol{w} + \boldsymbol{\xi}^\top \boldsymbol{Q}\,\boldsymbol{y}'(\overline{\boldsymbol{\xi}}) \\ = \; & \max_{\boldsymbol{\xi}\in\Xi} \; \boldsymbol{\xi}^\top \boldsymbol{C}\,\boldsymbol{x} + \boldsymbol{\xi}^\top \boldsymbol{D}\,\boldsymbol{w} + \boldsymbol{\xi}^\top \boldsymbol{Q}\,\boldsymbol{y}'(\boldsymbol{\xi}) \\ \le \; & \max_{\boldsymbol{\xi}\in\Xi} \; \boldsymbol{\xi}^\top \boldsymbol{C}\,\boldsymbol{x} + \boldsymbol{\xi}^\top \boldsymbol{D}\,\boldsymbol{w} + \boldsymbol{\xi}^\top \boldsymbol{Q}\,\boldsymbol{y}(\boldsymbol{\xi}). \end{aligned}$$

Thus, $(\boldsymbol{x},\boldsymbol{w})$ is feasible in Problem ($\mathcal{P}$) with a cost no greater than that of $(\boldsymbol{x},\boldsymbol{w},\boldsymbol{y}(\cdot))$ in Problem (3). We have thus shown that Problem ($\mathcal{P}$) lower bounds Problem (3) and that if $(\boldsymbol{x},\boldsymbol{w},\boldsymbol{y}(\cdot))$ is optimal in Problem (3), then $(\boldsymbol{x},\boldsymbol{w})$ is feasible in Problem ($\mathcal{P}$) with the cost attained by $\boldsymbol{x}$ in Problem ($\mathcal{P}$) being no greater than the cost of $(\boldsymbol{x},\boldsymbol{y}(\cdot))$ in Problem (3).

We conclude that the optimal costs of Problems (3) and ($\mathcal{P}$) are equal, and that claims *(i)* and *(ii)* hold. $\square$

Proof of Lemma 1 Fix $\boldsymbol{x} \in \mathcal{X}$, $\boldsymbol{w} \in \mathcal{W}$, and $\boldsymbol{y}^k$, $k \in \mathcal{K}$, and $\overline{\boldsymbol{\xi}} \in \Xi$. It suffices to show that the problems

$$\min_{k\in\mathcal{K}} \left\{ \max_{\boldsymbol{\xi}\in\Xi(\boldsymbol{w},\overline{\boldsymbol{\xi}})} \; \boldsymbol{\xi}^\top \boldsymbol{C}\,\boldsymbol{x} + \boldsymbol{\xi}^\top \boldsymbol{D}\,\boldsymbol{w} + \boldsymbol{\xi}^\top \boldsymbol{Q}\,\boldsymbol{y}^k \;:\; \boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \le \boldsymbol{H}\boldsymbol{\xi} \quad \forall \boldsymbol{\xi} \in \Xi(\boldsymbol{w},\overline{\boldsymbol{\xi}}) \right\} \tag{EC.5}$$

and

$$\max_{\boldsymbol{\xi}^k\in\Xi(\boldsymbol{w},\overline{\boldsymbol{\xi}}),\, k\in\mathcal{K}} \; \min_{k\in\mathcal{K}} \left\{ (\boldsymbol{\xi}^k)^\top \boldsymbol{C}\,\boldsymbol{x} + (\boldsymbol{\xi}^k)^\top \boldsymbol{D}\,\boldsymbol{w} + (\boldsymbol{\xi}^k)^\top \boldsymbol{Q}\,\boldsymbol{y}^k \;:\; \boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \le \boldsymbol{H}\boldsymbol{\xi}^k \right\} \tag{EC.6}$$

have the same optimal objective.

Problem (EC.5) is either infeasible or has a finite objective value. Indeed, it cannot be unbounded below since, if it is feasible, its objective value is given as the minimum of finitely many terms each of which is bounded, by virtue of the compactness of the non-empty set $\Xi(\boldsymbol{w},\overline{\boldsymbol{\xi}})$. Similarly, Problem (EC.6) is either unbounded above or has a finite objective value. It cannot be infeasible since $\Xi(\boldsymbol{w},\overline{\boldsymbol{\xi}})$ is non-empty.

We proceed in two steps. First, we show that Problem (EC.5) is infeasible if and only if Problem (EC.6) is unbounded above, in which case both problems have an optimal objective value of $+\infty$. Second, we show that if the problems have a finite optimal objective value, then their optimal values are equal.

For the first claim, we have

$$\begin{aligned}
&\text{Problem (EC.5) is infeasible}\\
\Leftrightarrow\ & \nexists k \in \mathcal{K} \,:\, \boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{H}\boldsymbol{\xi} \quad \forall \boldsymbol{\xi} \in \Xi(\boldsymbol{w}, \overline{\boldsymbol{\xi}})\\
\Leftrightarrow\ & \forall k \in \mathcal{K},\ \exists \tilde{\boldsymbol{\xi}}^k \in \Xi(\boldsymbol{w}, \overline{\boldsymbol{\xi}}) \,:\, \boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \not\leq \boldsymbol{H}\tilde{\boldsymbol{\xi}}^k\\
\Leftrightarrow\ & \text{Problem (EC.6) is unbounded.}
\end{aligned}$$

For the second claim, we proceed in two steps. First, we show that the optimal objective value of Problem (EC.6) can be no greater than the optimal objective value of Problem (EC.5). Then, we show that the converse is also true.

For the first part, let $\tilde{k}$ be feasible in Problem (EC.5) and $\{\tilde{\boldsymbol{\xi}}^k\}_{k\in\mathcal{K}}$ be feasible in Problem (EC.6). The objective value attained by $\tilde{k}$ in Problem (EC.5) is given by

$$\max_{\boldsymbol{\xi} \in \Xi(\boldsymbol{w}, \overline{\boldsymbol{\xi}})} \ \boldsymbol{\xi}^\top \boldsymbol{C}\,\boldsymbol{x} + \boldsymbol{\xi}^\top \boldsymbol{D}\,\boldsymbol{w} + \boldsymbol{\xi}^\top \boldsymbol{Q}\,\boldsymbol{y}^{\tilde{k}}.$$

Accordingly, the objective value attained by $\{\tilde{\boldsymbol{\xi}}^k\}_{k\in\mathcal{K}}$ in Problem (EC.6) is given by

$$\min_{k\in\mathcal{K}} \left\{ (\tilde{\boldsymbol{\xi}}^k)^\top \boldsymbol{C}\,\boldsymbol{x} + (\tilde{\boldsymbol{\xi}}^k)^\top \boldsymbol{D}\,\boldsymbol{w} + (\tilde{\boldsymbol{\xi}}^k)^\top \boldsymbol{Q}\,\boldsymbol{y}^k \,:\, \boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{H}\tilde{\boldsymbol{\xi}}^k \right\}.$$

Next, note that

$$\begin{aligned}
&\min_{k\in\mathcal{K}} \left\{ (\tilde{\boldsymbol{\xi}}^k)^\top \boldsymbol{C}\,\boldsymbol{x} + (\tilde{\boldsymbol{\xi}}^k)^\top \boldsymbol{D}\,\boldsymbol{w} + (\tilde{\boldsymbol{\xi}}^k)^\top \boldsymbol{Q}\,\boldsymbol{y}^k \,:\, \boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{H}\tilde{\boldsymbol{\xi}}^k \right\}\\
\leq\ & (\tilde{\boldsymbol{\xi}}^{\tilde{k}})^\top \boldsymbol{C}\,\boldsymbol{x} + (\tilde{\boldsymbol{\xi}}^{\tilde{k}})^\top \boldsymbol{D}\,\boldsymbol{w} + (\tilde{\boldsymbol{\xi}}^{\tilde{k}})^\top \boldsymbol{Q}\,\boldsymbol{y}^{\tilde{k}}\\
\leq\ & \max_{\boldsymbol{\xi} \in \Xi(\boldsymbol{w}, \overline{\boldsymbol{\xi}})} \ \boldsymbol{\xi}^\top \boldsymbol{C}\,\boldsymbol{x} + \boldsymbol{\xi}^\top \boldsymbol{D}\,\boldsymbol{w} + \boldsymbol{\xi}^\top \boldsymbol{Q}\,\boldsymbol{y}^{\tilde{k}},
\end{aligned}$$

where the first inequality follows by feasibility of $\tilde{k}$ in Problem (EC.5) since $\tilde{\boldsymbol{\xi}}^{\tilde{k}} \in \Xi(\boldsymbol{w}, \overline{\boldsymbol{\xi}})$ and the second inequality follows by feasibility of $\tilde{\boldsymbol{\xi}}^{\tilde{k}}$ in the maximization problem. Since the choices of $\tilde{k}$ and $\{\tilde{\boldsymbol{\xi}}^k\}_{k\in\mathcal{K}}$ were arbitrary, it follows that the optimal objective of Problem (EC.5) upper bounds the optimal objective of Problem (EC.6).

For the second part, we show that the converse also holds. For each $k \in \mathcal{K}$, let

$$\boldsymbol{\xi}^{k,\star} \in \operatorname*{arg\,max}_{\boldsymbol{\xi} \in \Xi(\boldsymbol{w}, \overline{\boldsymbol{\xi}})} \ \boldsymbol{\xi}^\top \boldsymbol{C}\,\boldsymbol{x} + \boldsymbol{\xi}^\top \boldsymbol{D}\,\boldsymbol{w} + \boldsymbol{\xi}^\top \boldsymbol{Q}\,\boldsymbol{y}^k.$$

Then, the optimal objective value of Problem (EC.5) is expressible as

$$\min_{k\in\mathcal{K}} \left\{ (\boldsymbol{\xi}^{k,\star})^\top \boldsymbol{C}\,\boldsymbol{x} + (\boldsymbol{\xi}^{k,\star})^\top \boldsymbol{D}\,\boldsymbol{w} + (\boldsymbol{\xi}^{k,\star})^\top \boldsymbol{Q}\,\boldsymbol{y}^k \;:\; \boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{H}\boldsymbol{\xi} \quad \forall \boldsymbol{\xi} \in \Xi(\boldsymbol{w},\overline{\boldsymbol{\xi}}) \right\}. \tag{EC.7}$$

Since $\boldsymbol{\xi}^{k,\star} \in \Xi(\boldsymbol{w},\overline{\boldsymbol{\xi}})$, the solution $\{\boldsymbol{\xi}^{k,\star}\}_{k\in\mathcal{K}}$ is feasible in Problem (EC.6) with objective

$$\min_{k\in\mathcal{K}} \left\{ (\boldsymbol{\xi}^{k,\star})^\top \boldsymbol{C}\,\boldsymbol{x} + (\boldsymbol{\xi}^{k,\star})^\top \boldsymbol{D}\,\boldsymbol{w} + (\boldsymbol{\xi}^{k,\star})^\top \boldsymbol{Q}\,\boldsymbol{y}^k \;:\; \boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{H}\boldsymbol{\xi}^{k,\star} \right\}. \tag{EC.8}$$

If the optimal objective values of Problems (EC.7) and (EC.8) are equal, then we can directly conclude that the optimal objective value of Problem (EC.6) exceeds that of Problem (EC.5). Suppose to the contrary that the optimal objective value of Problems (EC.8) is strictly lower than that of Problem (EC.7). Then, there exists (at least one) $k \in \mathcal{K}$ that is feasible in (EC.8) but infeasible in (EC.7) and for each such k, there exists $\boldsymbol{\xi}^{k,\prime} \in \Xi(\boldsymbol{w},\overline{\boldsymbol{\xi}})$ such that $\boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^{\tilde{k}} \not\leq \boldsymbol{H}\boldsymbol{\xi}^{k,\prime}$. We can construct a feasible solution $\{\tilde{\boldsymbol{\xi}}^{k,\star}\}_{k\in\mathcal{K}}$ to Problem (EC.6) with the same objective as Problem (EC.7) as follows:

$$\tilde{\boldsymbol{\xi}}^{k,\star} := \begin{cases} \boldsymbol{\xi}^{k,\star} & \text{if } k \,:\, \boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{H}\boldsymbol{\xi} \quad \forall \boldsymbol{\xi} \in \Xi(\boldsymbol{w},\overline{\boldsymbol{\xi}}), \\[2ex] \boldsymbol{\xi}^{k,\prime} & \text{else.} \end{cases}$$

Indeed, the objective value attained by $\{\tilde{\boldsymbol{\xi}}^{k,\star}\}_{k\in\mathcal{K}}$ in Problem (EC.6) is

$$\begin{aligned} &\min_{k\in\mathcal{K}} \left\{ (\tilde{\boldsymbol{\xi}}^{k,\star})^\top \boldsymbol{C}\,\boldsymbol{x} + (\tilde{\boldsymbol{\xi}}^{k,\star})^\top \boldsymbol{D}\,\boldsymbol{w} + (\tilde{\boldsymbol{\xi}}^{k,\star})^\top \boldsymbol{Q}\,\boldsymbol{y}^k \;:\; \boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{H}\tilde{\boldsymbol{\xi}}^{k,\star} \right\} \\ = &\min_{k\in\mathcal{K}} \left\{ (\tilde{\boldsymbol{\xi}}^{k,\star})^\top \boldsymbol{C}\,\boldsymbol{x} + (\tilde{\boldsymbol{\xi}}^{k,\star})^\top \boldsymbol{D}\,\boldsymbol{w} + (\tilde{\boldsymbol{\xi}}^{k,\star})^\top \boldsymbol{Q}\,\boldsymbol{y}^k \;:\; \boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{H}\boldsymbol{\xi} \quad \forall \boldsymbol{\xi} \in \Xi(\boldsymbol{w},\overline{\boldsymbol{\xi}}) \right\} \\ = &\min_{k\in\mathcal{K}} \left\{ (\boldsymbol{\xi}^{k,\star})^\top \boldsymbol{C}\,\boldsymbol{x} + (\boldsymbol{\xi}^{k,\star})^\top \boldsymbol{D}\,\boldsymbol{w} + (\boldsymbol{\xi}^{k,\star})^\top \boldsymbol{Q}\,\boldsymbol{y}^k \;:\; \boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{H}\boldsymbol{\xi} \quad \forall \boldsymbol{\xi} \in \Xi(\boldsymbol{w},\overline{\boldsymbol{\xi}}) \right\}, \end{aligned}$$

where the first equality follows by construction since

$$\{k \in \mathcal{K} \,:\, \boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{H}\boldsymbol{\xi} \quad \forall \boldsymbol{\xi} \in \Xi(\boldsymbol{w},\overline{\boldsymbol{\xi}})\} = \{k \in \mathcal{K} \,:\, \boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{H}\tilde{\boldsymbol{\xi}}^{k,\star}\}$$

and the second equality follows since

$$\tilde{\boldsymbol{\xi}}^{k,\star} = \boldsymbol{\xi}^{k,\star} \quad \forall k \in \mathcal{K} \,:\, \boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{H}\boldsymbol{\xi} \quad \forall \boldsymbol{\xi} \in \Xi(\boldsymbol{w},\overline{\boldsymbol{\xi}}).$$

We have thus shown that the optimal objective value of Problem (EC.6) is at least as large as that of Problem (EC.5).

Combining the first and second parts of the proof, we conclude that Problems (EC.5) and (EC.6) have the same optimal objective values, which concludes the proof. $\square$

EC.3. Proofs of Statements in Section 4

Proof of Lemma 2 Since Problem ($\mathcal{PO}_K$) is equivalent to Problem (9) (by Lemma 1), it suffices to show that Problems (9) and (10) are equivalent.

First, note that for any choice of $\boldsymbol{w} \in \mathcal{W}$, the set $\Xi^K(\boldsymbol{w})$ is non-empty. If there is no $\boldsymbol{x} \in \mathcal{X}$, $\boldsymbol{w} \in \mathcal{W}$, and $\boldsymbol{y} \in \mathcal{Y}$ such that $\boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y} \leq \boldsymbol{h}$, then Problem (10) is infeasible and has an optimal objective value of $+\infty$. Accordingly, Problem (9) also has an objective value of $+\infty$ since either its outer or inner minimization problems are infeasible.

Suppose now that there exists $\boldsymbol{x} \in \mathcal{X}$, $\boldsymbol{w} \in \mathcal{W}$, and $\boldsymbol{y} \in \mathcal{Y}$ such that $\boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y} \leq \boldsymbol{h}$. Then, Problems (9) and (10) are both feasible. Let $(\boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}\}_{k\in\mathcal{K}})$ be a feasible solution for (10). Then, it is feasible in (9) and attains the same objective value in both problems since all second stage policies $\boldsymbol{y}^k$, $k \in \mathcal{K}$, satisfy the second-stage constraints in Problem (9). Conversely, let $(\boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}\}_{k\in\mathcal{K}})$ be feasible in Problem (9). Since $\Xi^K(\boldsymbol{w})$ is non-empty, there must exist $k^\star \in \mathcal{K}$ such that $\boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^{k^\star} \leq \boldsymbol{h}$ (else the problem would have an optimal objective value of $+\infty$ and thus be infeasible, a contradiction). If $\boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{h}$ for all $k \in \mathcal{K}$, then $(\boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}\}_{k\in\mathcal{K}})$ is feasible in (10) and attains the same objective value in both problems. On the other hand, if $\boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k > \boldsymbol{h}$ for some $k \in \mathcal{K}$, define

$$\overline{\boldsymbol{y}}^k = \begin{cases} \boldsymbol{y}^k & \text{if } \boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{h} \\ \boldsymbol{y}^{k^\star} & \text{else.} \end{cases}$$

Then, $(\boldsymbol{x}, \boldsymbol{w}, \{\overline{\boldsymbol{y}}\}_{k\in\mathcal{K}})$ is feasible in (10) and attains the same objective value in both problems. $\square$

Proof of Observation 1 Fix $K \in \mathbb{N}$ and $(\boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}^k\}_{k\in\mathcal{K}})$ such that $\boldsymbol{x} \in \mathcal{X}$, $\boldsymbol{w} \in \mathcal{W}$, $\boldsymbol{y}^k \in \mathcal{Y}$. Assume, w.l.o.g. (see the Proof of Lemma 2) that $\boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{h}$ for all $k \in \mathcal{K}$. From Lemma 2, the objective value of ($\mathcal{PO}_K$) under this decision is equal to

$$\begin{array}{ll} \text{maximize} & \min_{k\in\mathcal{K}} \left\{ (\boldsymbol{\xi}^k)^\top \boldsymbol{C}\,\boldsymbol{x} + (\boldsymbol{\xi}^k)^\top \boldsymbol{D}\,\boldsymbol{w} + (\boldsymbol{\xi}^k)^\top \boldsymbol{Q}\,\boldsymbol{y}^k \right\} \\ \text{subject to} & \overline{\boldsymbol{\xi}} \in \Xi,\ \boldsymbol{\xi}^k \in \Xi,\ k \in \mathcal{K} \\ & \boldsymbol{w} \circ \boldsymbol{\xi}^k = \boldsymbol{w} \circ \overline{\boldsymbol{\xi}} \quad \forall k \in \mathcal{K}. \end{array}$$

We can write the problem above in epigraph form as an LP:

$$\begin{array}{ll} \text{maximize} & \tau \\ \text{subject to} & \tau \in \mathbb{R},\ \overline{\boldsymbol{\xi}} \in \Xi,\ \boldsymbol{\xi}^k \in \Xi, k \in \mathcal{K} \\ & \tau \leq (\boldsymbol{\xi}^k)^\top \boldsymbol{C}\,\boldsymbol{x} + (\boldsymbol{\xi}^k)^\top \boldsymbol{D}\,\boldsymbol{w} + (\boldsymbol{\xi}^k)^\top \boldsymbol{Q}\,\boldsymbol{y}^k \quad \forall k \in \mathcal{K} \\ & \boldsymbol{w} \circ \boldsymbol{\xi}^k = \boldsymbol{w} \circ \overline{\boldsymbol{\xi}} \quad \forall k \in \mathcal{K}. \end{array}$$

For any fixed K, the size of this LP is polynomial in the size of the input. $\square$

Proof of Theorem 3 For any fixed $(\boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}^k\}_{k\in\mathcal{K}})$, we can express the inner maximization problem in (10) in epigraph form as

$$
\begin{array}{ll}
\text{maximize} & \tau \\
\text{subject to} & \tau \in \mathbb{R},\, \overline{\boldsymbol{\xi}} \in \mathbb{R}^{N_\xi},\, \boldsymbol{\xi}^k \in \mathbb{R}^{N_\xi}, k \in \mathcal{K} \\
& \tau \leq (\boldsymbol{C}\,\boldsymbol{x} + \boldsymbol{D}\,\boldsymbol{w} + \boldsymbol{Q}\,\boldsymbol{y}^k)^\top \boldsymbol{\xi}^k \quad \forall k \in \mathcal{K} \\
& \boldsymbol{A}\overline{\boldsymbol{\xi}} \leq \boldsymbol{b} \\
& \boldsymbol{A}\boldsymbol{\xi}^k \leq \boldsymbol{b} \quad \forall k \in \mathcal{K} \\
& \boldsymbol{w} \circ \boldsymbol{\xi}^k = \boldsymbol{w} \circ \overline{\boldsymbol{\xi}} \quad \forall k \in \mathcal{K}.
\end{array}
$$

Strong LP duality (which applies since the feasible set is non-empty and since the problem is bounded by virtue of the boundedness of Ξ) implies that the optimal objective value of this problem coincides with the optimal objective value of its dual

$$
\begin{array}{ll}
\text{minimize} & \boldsymbol{b}^\top \boldsymbol{\beta} + \displaystyle\sum_{k\in\mathcal{K}} \boldsymbol{b}^\top \boldsymbol{\beta}^k \\
\text{subject to} & \boldsymbol{\alpha} \in \mathbb{R}^K_+,\, \boldsymbol{\beta} \in \mathbb{R}^R_+,\, \boldsymbol{\beta}^k \in \mathbb{R}^R_+,\, \boldsymbol{\gamma}^k \in \mathbb{R}^{N_\xi},\, k \in \mathcal{K} \\
& \mathbf{e}^\top \boldsymbol{\alpha} = 1 \\
& \boldsymbol{A}^\top \boldsymbol{\beta}^k + \boldsymbol{w} \circ \boldsymbol{\gamma}^k = \boldsymbol{\alpha}_k\,(\boldsymbol{C}\boldsymbol{x} + \boldsymbol{D}\boldsymbol{w} + \boldsymbol{Q}\boldsymbol{y}^k) \quad \forall k \in \mathcal{K} \\
& \boldsymbol{A}^\top \boldsymbol{\beta} = \displaystyle\sum_{k\in\mathcal{K}} \boldsymbol{w} \circ \boldsymbol{\gamma}^k.
\end{array}
$$

We can now group the outer minimization with the minimization above to obtain

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

The result now follows by replacing the bilinear terms $\boldsymbol{w} \circ \boldsymbol{\gamma}^k$, $\boldsymbol{\alpha}_k \boldsymbol{x}$, $\boldsymbol{\alpha}_k \boldsymbol{w}$, and $\boldsymbol{\alpha}_k \boldsymbol{y}^k$ with auxiliary variables $\overline{\boldsymbol{\gamma}}^k \in \mathbb{R}^{N_\xi}$, $\overline{\boldsymbol{x}}^k \in \mathbb{R}_+^{N_x}$, $\overline{\boldsymbol{w}}^k \in \mathbb{R}_+^{N_\xi}$, and $\overline{\boldsymbol{y}}^k \in \mathbb{R}_+^{N_y}$ such that

$$
\begin{aligned}
\overline{\boldsymbol{\gamma}}^k = \boldsymbol{w} \circ \boldsymbol{\gamma}^k \;&\Leftrightarrow\; \overline{\boldsymbol{\gamma}}^k \leq \boldsymbol{\gamma}^k + M(\mathbf{e} - \boldsymbol{w}),\, \overline{\boldsymbol{\gamma}}^k \leq M\boldsymbol{w},\, \overline{\boldsymbol{\gamma}}^k \geq -M\boldsymbol{w},\, \overline{\boldsymbol{\gamma}}^k \geq \boldsymbol{\gamma}^k - M(\mathbf{e} - \boldsymbol{w}),\\
\overline{\boldsymbol{x}}^k = \boldsymbol{\alpha}_k \boldsymbol{x} \;&\Leftrightarrow\; \overline{\boldsymbol{x}}^k \leq \boldsymbol{x},\, \overline{\boldsymbol{x}}^k \leq \boldsymbol{\alpha}_k \mathbf{e},\, \overline{\boldsymbol{x}}^k \geq (\boldsymbol{\alpha}_k - 1)\mathbf{e} + \boldsymbol{x},\\
\overline{\boldsymbol{w}}^k = \boldsymbol{\alpha}_k \boldsymbol{w} \;&\Leftrightarrow\; \overline{\boldsymbol{w}}^k \leq \boldsymbol{w},\, \overline{\boldsymbol{w}}^k \leq \boldsymbol{\alpha}_k \mathbf{e},\, \overline{\boldsymbol{w}}^k \geq (\boldsymbol{\alpha}_k - 1)\mathbf{e} + \boldsymbol{w},\\
\overline{\boldsymbol{y}}^k = \boldsymbol{\alpha}_k \boldsymbol{y}^k \;&\Leftrightarrow\; \overline{\boldsymbol{y}}^k \leq \boldsymbol{y}^k,\, \overline{\boldsymbol{y}}^k \leq \boldsymbol{\alpha}_k \mathbf{e},\, \overline{\boldsymbol{y}}^k \geq (\boldsymbol{\alpha}_k - 1)\mathbf{e} + \boldsymbol{y}^k,
\end{aligned}
$$

where in the last three cases we have exploited the fact that $\boldsymbol{x}$, $\boldsymbol{w}$, and $\boldsymbol{y}^k$ are binary and that $\boldsymbol{\alpha}^k \in [\mathbf{0}, \mathbf{e}]$. □

EC.4. Proofs of Statements in Section 5

Proof of Theorem 4 The proof is a direct consequence of Theorem 3 in Hanasusanto et al. (2015). Indeed, the authors show that evaluating the objective function of Problem (2) is strongly NP-hard. Since Problem (2) can be reduced in polynomial time to an instance of Problem ($\mathcal{P}$) by letting $\boldsymbol{D} = \mathbf{0}$, $\boldsymbol{V} = \mathbf{0}$, and $\boldsymbol{w} = \mathbf{e}$, this concludes the proof. □

The proof below is a generalization of the proof of Proposition 1 in Hanasusanto et al. (2015) that operates in the *lifted* uncertainty and decision spaces. Despite this key difference, the proof idea carries through.

Proof of Proposition 1 Fix $\boldsymbol{x}$, $\boldsymbol{w}$, and $\{\boldsymbol{y}^k\}_{k\in\mathcal{K}}$. We show that $\{\Xi^K(\boldsymbol{w}, \boldsymbol{\ell})\}_{\boldsymbol{\ell}\in\mathcal{L}}$ is a cover of $\Xi^K(\boldsymbol{w})$, i.e., that $\Xi^K(\boldsymbol{w}) = \bigcup_{\boldsymbol{\ell}\in\mathcal{L}} \Xi^K(\boldsymbol{w}, \boldsymbol{\ell})$. Let $\{\boldsymbol{\xi}^k\}_{k\in\mathcal{K}} \in \Xi^K(\boldsymbol{w})$ and define

$$
\boldsymbol{\ell}_k = \begin{cases} 0, \text{ if } \boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{H}\boldsymbol{\xi}^k \\ \\ \min\left\{\ell \in \{1, \ldots, L\} \,:\, [\boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k]_\ell > [\boldsymbol{H}\boldsymbol{\xi}^k]_\ell\right\}, \text{ else.} \end{cases} \qquad \forall k \in \mathcal{K}.
$$

Then, $\{\boldsymbol{\xi}^k\}_{k\in\mathcal{K}} \in \Xi^K(\boldsymbol{w}, \boldsymbol{\ell})$. Moreover, by definition, we have $\Xi^K(\boldsymbol{w}, \boldsymbol{\ell}) \subseteq \Xi^K(\boldsymbol{w})$ for all $\boldsymbol{\ell} \in \mathcal{L}$. Therefore $\{\Xi^K(\boldsymbol{w}, \boldsymbol{\ell})\}_{\boldsymbol{\ell}\in\mathcal{L}}$ is a cover of $\Xi^K(\boldsymbol{w})$. It then follows that

$$
\begin{aligned}
&\max_{\{\boldsymbol{\xi}^k\}_{k\in\mathcal{K}} \in \Xi^K(\boldsymbol{w})} \; \min_{k\in\mathcal{K}} \left\{(\boldsymbol{\xi}^k)^\top \boldsymbol{C}\,\boldsymbol{x} + (\boldsymbol{\xi}^k)^\top \boldsymbol{D}\,\boldsymbol{w} + (\boldsymbol{\xi}^k)^\top \boldsymbol{Q}\,\boldsymbol{y}^k \,:\, \boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{H}\boldsymbol{\xi}^k\right\}\\
=\; &\max_{\{\boldsymbol{\xi}^k\}_{k\in\mathcal{K}} \in \bigcup_{\boldsymbol{\ell}\in\mathcal{L}} \Xi^K(\boldsymbol{w}, \boldsymbol{\ell})} \; \min_{k\in\mathcal{K}} \left\{(\boldsymbol{\xi}^k)^\top \boldsymbol{C}\,\boldsymbol{x} + (\boldsymbol{\xi}^k)^\top \boldsymbol{D}\,\boldsymbol{w} + (\boldsymbol{\xi}^k)^\top \boldsymbol{Q}\,\boldsymbol{y}^k \,:\, \boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{H}\boldsymbol{\xi}^k\right\}\\
=\; &\max_{\boldsymbol{\ell}\in\mathcal{L}} \; \max_{\{\boldsymbol{\xi}^k\}_{k\in\mathcal{K}} \in \Xi^K(\boldsymbol{w}, \boldsymbol{\ell})} \; \min_{k\in\mathcal{K}} \left\{(\boldsymbol{\xi}^k)^\top \boldsymbol{C}\,\boldsymbol{x} + (\boldsymbol{\xi}^k)^\top \boldsymbol{D}\,\boldsymbol{w} + (\boldsymbol{\xi}^k)^\top \boldsymbol{Q}\,\boldsymbol{y}^k \,:\, \boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{H}\boldsymbol{\xi}^k\right\}.
\end{aligned}
$$

The definition of $\Xi^K(\boldsymbol{w}, \boldsymbol{\ell})$ implies that $\boldsymbol{\ell}_k = 0$ if and only if $\boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{H}\boldsymbol{\xi}^k$. This concludes the proof. □

Proof of Theorem 5 The objective function of the approximate problem (12_ϵ) is identical to

$$
\max_{\boldsymbol{\ell}\in\mathcal{L}} \; \max_{\{\boldsymbol{\xi}^k\}_{k\in\mathcal{K}} \in \Xi_\epsilon^K(\boldsymbol{w}, \boldsymbol{\ell})} \; \min_{\boldsymbol{\lambda}\in\Lambda_K(\boldsymbol{\ell})} \left\{\sum_{k\in\mathcal{K}} \boldsymbol{\lambda}_k \left[(\boldsymbol{\xi}^k)^\top \boldsymbol{C}\,\boldsymbol{x} + (\boldsymbol{\xi}^k)^\top \boldsymbol{D}\,\boldsymbol{w} + (\boldsymbol{\xi}^k)^\top \boldsymbol{Q}\,\boldsymbol{y}^k\right]\right\},
$$

where $\Lambda_K(\boldsymbol{\ell}) := \left\{\boldsymbol{\lambda} \in \mathbb{R}^K_+ : \mathbf{e}^\top \boldsymbol{\lambda} = 1,\, \boldsymbol{\lambda}_k = 0 \,\forall k \in \mathcal{K} : \boldsymbol{\ell}_k \neq 0\right\}$. Note that $\Lambda_K(\boldsymbol{\ell}) = \emptyset$ if and only if $\boldsymbol{\ell} > \mathbf{0}$. If $\Xi^K_\epsilon(\boldsymbol{w}, \boldsymbol{\ell}) = \emptyset$ for all $\boldsymbol{\ell} \in \mathcal{L}_+$, then the problem is equivalent to

$$\max_{\boldsymbol{\ell} \in \partial \mathcal{L}} \; \max_{\{\boldsymbol{\xi}^k\}_{k \in \mathcal{K}} \in \Xi^K_\epsilon(\boldsymbol{w}, \boldsymbol{\ell})} \; \min_{\boldsymbol{\lambda} \in \Lambda_K(\boldsymbol{\ell})} \left\{ \sum_{k \in \mathcal{K}} \boldsymbol{\lambda}_k \left[(\boldsymbol{\xi}^k)^\top \boldsymbol{C}\, \boldsymbol{x} + (\boldsymbol{\xi}^k)^\top \boldsymbol{D}\, \boldsymbol{w} + (\boldsymbol{\xi}^k)^\top \boldsymbol{Q}\, \boldsymbol{y}^k \right] \right\},$$

and we can apply the classical min-max theorem (since $\Lambda_K(\boldsymbol{\ell})$ is nonempty for all $\boldsymbol{\ell} \in \partial\mathcal{L}$) to obtain the equivalent reformulation

$$\max_{\boldsymbol{\ell} \in \partial \mathcal{L}} \; \min_{\boldsymbol{\lambda} \in \Lambda_K(\boldsymbol{\ell})} \; \max_{\{\boldsymbol{\xi}^k\}_{k \in \mathcal{K}} \in \Xi^K_\epsilon(\boldsymbol{w}, \boldsymbol{\ell})} \left\{ \sum_{k \in \mathcal{K}} \boldsymbol{\lambda}_k \left[(\boldsymbol{\xi}^k)^\top \boldsymbol{C}\, \boldsymbol{x} + (\boldsymbol{\xi}^k)^\top \boldsymbol{D}\, \boldsymbol{w} + (\boldsymbol{\xi}^k)^\top \boldsymbol{Q}\, \boldsymbol{y}^k \right] \right\},$$

which in turn is equivalent to

$$\min_{\substack{\boldsymbol{\lambda}(\boldsymbol{\ell}) \in \Lambda_K(\boldsymbol{\ell}), \\ \boldsymbol{\ell} \in \partial \mathcal{L}}} \; \max_{\boldsymbol{\ell} \in \partial \mathcal{L}} \; \max_{\{\boldsymbol{\xi}^k\}_{k \in \mathcal{K}} \in \Xi^K_\epsilon(\boldsymbol{w}, \boldsymbol{\ell})} \left\{ \sum_{k \in \mathcal{K}} \boldsymbol{\lambda}_k(\boldsymbol{\ell}) \left[(\boldsymbol{\xi}^k)^\top \boldsymbol{C}\, \boldsymbol{x} + (\boldsymbol{\xi}^k)^\top \boldsymbol{D}\, \boldsymbol{w} + (\boldsymbol{\xi}^k)^\top \boldsymbol{Q}\, \boldsymbol{y}^k \right] \right\}.$$

If, on the other hand, $\Xi^K_\epsilon(\boldsymbol{w}, \boldsymbol{\ell}) \neq \emptyset$ for some $\boldsymbol{\ell} \in \mathcal{L}_+$, then the objective function in (12_ϵ) evaluates to $+\infty$. Using an epigraph reformulation, we thus conclude that (12_ϵ) is equivalent to the problem

$$\begin{array}{lll}
\text{minimize} & \tau & \\
\text{subject to} & \boldsymbol{x} \in \mathcal{X},\, \boldsymbol{w} \in \mathcal{W},\, \boldsymbol{y}^k \in \mathcal{Y},\, k \in \mathcal{K} & \\
& \tau \in \mathbb{R},\, \boldsymbol{\lambda}(\boldsymbol{\ell}) \in \Lambda_K(\boldsymbol{\ell}),\, \boldsymbol{\ell} \in \partial \mathcal{L} & \\
& \tau \geq \displaystyle\sum_{k \in \mathcal{K}} \boldsymbol{\lambda}_k(\boldsymbol{\ell}) \left[(\boldsymbol{\xi}^k)^\top \boldsymbol{C}\, \boldsymbol{x} + (\boldsymbol{\xi}^k)^\top \boldsymbol{D}\, \boldsymbol{w} + (\boldsymbol{\xi}^k)^\top \boldsymbol{Q}\, \boldsymbol{y}^k \right] & \forall \boldsymbol{\ell} \in \partial\mathcal{L},\, \{\boldsymbol{\xi}^k\}_{k \in \mathcal{K}} \in \Xi^K_\epsilon(\boldsymbol{w}, \boldsymbol{\ell}) \\
& \Xi^K_\epsilon(\boldsymbol{w}, \boldsymbol{\ell}) = \emptyset & \forall \boldsymbol{\ell} \in \mathcal{L}_+.
\end{array} \tag{EC.9}$$

The semi-infinite constraint associated with $\boldsymbol{\ell} \in \partial\mathcal{L}$ is satisfied if and only if the optimal value of

$$\begin{array}{lll}
\text{maximize} & \displaystyle\sum_{k \in \mathcal{K}} \boldsymbol{\lambda}_k(\boldsymbol{\ell}) \left[(\boldsymbol{\xi}^k)^\top \boldsymbol{C}\, \boldsymbol{x} + (\boldsymbol{\xi}^k)^\top \boldsymbol{D}\, \boldsymbol{w} + (\boldsymbol{\xi}^k)^\top \boldsymbol{Q}\, \boldsymbol{y}^k \right] & \\
\text{subject to} & \overline{\boldsymbol{\xi}} \in \mathbb{R}^{N_\xi},\, \boldsymbol{\xi}^k \in \mathbb{R}^{N_\xi},\, k \in \mathcal{K} & \\
& \boldsymbol{A}\overline{\boldsymbol{\xi}} \leq \boldsymbol{b} & \\
& \boldsymbol{A}\boldsymbol{\xi}^k \leq \boldsymbol{b} & \forall k \in \mathcal{K} \\
& \boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{H}\boldsymbol{\xi}^k & \forall k \in \mathcal{K} : \boldsymbol{\ell}_k = 0 \\
& [\boldsymbol{T}\boldsymbol{x} + \boldsymbol{V}\boldsymbol{w} + \boldsymbol{W}\boldsymbol{y}^k]_{\boldsymbol{\ell}_k} \geq [\boldsymbol{H}\boldsymbol{\xi}^k]_{\boldsymbol{\ell}_k} + \epsilon & \forall k \in \mathcal{K} : \boldsymbol{\ell}_k \neq 0 \\
& \boldsymbol{w} \circ \boldsymbol{\xi}^k = \boldsymbol{w} \circ \overline{\boldsymbol{\xi}} & \forall k \in \mathcal{K}
\end{array}$$

does not exceed τ. Strong linear programming duality implies that this problem attains the same optimal value as its dual problem which is given by

$$\begin{array}{ll}
\text{minimize} & \boldsymbol{b}^\top\left(\boldsymbol{\alpha}+\sum_{k\in\mathcal{K}}\boldsymbol{\alpha}^k\right)-\sum_{\substack{k\in\mathcal{K}:\\ \boldsymbol{\ell}_k=0}}(\boldsymbol{T}\boldsymbol{x}+\boldsymbol{V}\boldsymbol{w}+\boldsymbol{W}\boldsymbol{y}^k)^\top\boldsymbol{\beta}^k+\sum_{\substack{k\in\mathcal{K}:\\ \boldsymbol{\ell}_k\neq 0}}\left(\left[\boldsymbol{T}\boldsymbol{x}+\boldsymbol{V}\boldsymbol{w}+\boldsymbol{W}\boldsymbol{y}^k\right]_{\boldsymbol{\ell}_k}-\epsilon\right)\boldsymbol{\gamma}_k \\
\text{subject to} & \boldsymbol{\alpha}\in\mathbb{R}^R_+,\ \boldsymbol{\alpha}^k\in\mathbb{R}^R_+,\ \boldsymbol{\beta}^k\in\mathbb{R}^L_+,\ k\in\mathcal{K},\ \boldsymbol{\gamma}\in\mathbb{R}^K_+,\ \boldsymbol{\eta}^k\in\mathbb{R}^{N_\xi},\ k\in\mathcal{K} \\
& \boldsymbol{A}^\top\boldsymbol{\alpha}=\sum_{k\in\mathcal{K}}\boldsymbol{w}\circ\boldsymbol{\eta}^k \\
& \boldsymbol{A}^\top\boldsymbol{\alpha}^k-\boldsymbol{H}^\top\boldsymbol{\beta}^k+\boldsymbol{w}\circ\boldsymbol{\eta}^k=\boldsymbol{\lambda}_k(\boldsymbol{\ell})\left[\boldsymbol{C}\,\boldsymbol{x}+\boldsymbol{D}\,\boldsymbol{w}+\boldsymbol{Q}\,\boldsymbol{y}^k\right] \qquad \forall k\in\mathcal{K}:\boldsymbol{\ell}_k=0 \\
& \boldsymbol{A}^\top\boldsymbol{\alpha}^k+[\boldsymbol{H}]_{\boldsymbol{\ell}_k}\gamma_k+\boldsymbol{w}\circ\boldsymbol{\eta}^k=\boldsymbol{\lambda}_k(\boldsymbol{\ell})\left[\boldsymbol{C}\,\boldsymbol{x}+\boldsymbol{D}\,\boldsymbol{w}+\boldsymbol{Q}\,\boldsymbol{y}^k\right] \qquad \forall k\in\mathcal{K}:\boldsymbol{\ell}_k\neq 0.
\end{array}$$

Strong duality holds because the dual problem is always feasible. Indeed, one can show that the compactness of Ξ implies that $\{\boldsymbol{A}^\top\boldsymbol{\alpha}:\boldsymbol{\alpha}\geq\boldsymbol{0}\}=\mathbb{R}^{N_\xi}$. Note that the first constraint set in Problem (13) ensures that the optimal value of this dual problem does not exceed τ for all $\boldsymbol{\ell}\in\partial\mathcal{L}$.

The last constraint in (EC.9) is satisfied for $\boldsymbol{\ell}\in\mathcal{L}_+$ whenever the linear program

$$\begin{array}{lll}
\text{maximize} & 0 & \\
\text{subject to} & \overline{\boldsymbol{\xi}}\in\mathbb{R}^{N_\xi},\ \boldsymbol{\xi}^k\in\mathbb{R}^{N_\xi},\ k\in\mathcal{K} & \\
& \boldsymbol{A}\overline{\boldsymbol{\xi}}\leq\boldsymbol{b} & \\
& \boldsymbol{A}\boldsymbol{\xi}^k\leq\boldsymbol{b} & \forall k\in\mathcal{K} \\
& [\boldsymbol{T}\boldsymbol{x}+\boldsymbol{V}\boldsymbol{w}+\boldsymbol{W}\boldsymbol{y}^k]_{\boldsymbol{\ell}_k}\geq[\boldsymbol{H}\boldsymbol{\xi}^k]_{\boldsymbol{\ell}_k}+\epsilon & \forall k\in\mathcal{K} \\
& \boldsymbol{w}\circ\boldsymbol{\xi}^k=\boldsymbol{w}\circ\overline{\boldsymbol{\xi}} & \forall k\in\mathcal{K}
\end{array}$$

is infeasible. The dual to this problem reads

$$\begin{array}{ll}
\text{minimize} & \boldsymbol{b}^\top\left(\boldsymbol{\alpha}+\sum_{k\in\mathcal{K}}\boldsymbol{\alpha}^k\right)+\sum_{k\in\mathcal{K}}\left(\left[\boldsymbol{T}\boldsymbol{x}+\boldsymbol{V}\boldsymbol{w}+\boldsymbol{W}\boldsymbol{y}^k\right]_{\boldsymbol{\ell}_k}-\epsilon\right)\gamma_k \\
\text{subject to} & \boldsymbol{\alpha}\in\mathbb{R}^R_+,\ \boldsymbol{\alpha}^k\in\mathbb{R}^R_+,\ k\in\mathcal{K},\ \boldsymbol{\gamma}\in\mathbb{R}^K_+,\ \boldsymbol{\eta}^k\in\mathbb{R}^{N_\xi},\ k\in\mathcal{K} \\
& \boldsymbol{A}^\top\boldsymbol{\alpha}=\sum_{k\in\mathcal{K}}\boldsymbol{w}\circ\boldsymbol{\eta}^k \\
& \boldsymbol{A}^\top\boldsymbol{\alpha}^k+[\boldsymbol{H}]_{\boldsymbol{\ell}_k}\gamma_k+\boldsymbol{w}\circ\boldsymbol{\eta}^k=\boldsymbol{0} \qquad \forall k\in\mathcal{K}.
\end{array}$$

The feasible set of this dual is a cone and thus feasible (set $\boldsymbol{\alpha}=\boldsymbol{0}$, $\boldsymbol{\eta}^k=\boldsymbol{0}$, $\boldsymbol{\gamma}_k=0$, $k\in\mathcal{K}$). Therefore, strong LP duality applies and the primal is infeasible if and only if the dual is unbounded. Since the feasible set of the dual is a cone, the dual is unbounded if and only if there exists a feasible solution attaining an objective value of -1. $\square$

Proof of Observation 2 Suppose that we are only in the presence of exogenous uncertainty, i.e., $\boldsymbol{w} = \mathbf{e}$, $\boldsymbol{D} = \mathbf{0}$, and $\boldsymbol{V} = \mathbf{0}$. Then, Problem (13) reduces to

$$
\begin{array}{ll}
\min & \tau \\
\text{s.t.} & \tau \in \mathbb{R},\ \boldsymbol{x} \in \mathcal{X},\ \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},\ \boldsymbol{\gamma}(\boldsymbol{\ell}) \in \mathbb{R}^K_+,\ \boldsymbol{\eta}^k(\boldsymbol{\ell}) \in \mathbb{R}^{N_\xi},\ k \in \mathcal{K}, \boldsymbol{\ell} \in \mathcal{L} \\
& \left.\begin{array}{l}
\boldsymbol{\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}) = \displaystyle\sum_{k \in \mathcal{K}} \boldsymbol{\eta}^k(\boldsymbol{\ell}) \\
\boldsymbol{A}^\top \boldsymbol{\alpha}^k(\boldsymbol{\ell}) - \boldsymbol{H}^\top \boldsymbol{\beta}^k(\boldsymbol{\ell}) + \boldsymbol{\eta}^k(\boldsymbol{\ell}) = \boldsymbol{\lambda}_k(\boldsymbol{\ell}) \, [\boldsymbol{C}\,\boldsymbol{x} + \boldsymbol{Q}\,\boldsymbol{y}^k] \quad \forall k \in \mathcal{K} : \boldsymbol{\ell}_k = 0 \\
\boldsymbol{A}^\top \boldsymbol{\alpha}^k(\boldsymbol{\ell}) + [\boldsymbol{H}]_{\boldsymbol{\ell}_k} \gamma_k(\boldsymbol{\ell}) + \boldsymbol{\eta}^k(\boldsymbol{\ell}) = \boldsymbol{\lambda}_k(\boldsymbol{\ell}) \, [\boldsymbol{C}\,\boldsymbol{x} + \boldsymbol{Q}\,\boldsymbol{y}^k] \quad \forall k \in \mathcal{K} : \boldsymbol{\ell}_k \neq 0 \\
\tau \geq \boldsymbol{b}^\top \left(\boldsymbol{\alpha}(\boldsymbol{\ell}) + \displaystyle\sum_{k \in \mathcal{K}} \boldsymbol{\alpha}^k(\boldsymbol{\ell}) \right) - \displaystyle\sum_{\substack{k \in \mathcal{K}: \\ \boldsymbol{\ell}_k = 0}} (\boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^k)^\top \boldsymbol{\beta}^k(\boldsymbol{\ell}) \\
\qquad\qquad + \displaystyle\sum_{\substack{k \in \mathcal{K}: \\ \boldsymbol{\ell}_k \neq 0}} \left(\left[\boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^k \right]_{\boldsymbol{\ell}_k} - \epsilon \right) \gamma_k(\boldsymbol{\ell})
\end{array}\right\} \quad \forall \boldsymbol{\ell} \in \partial\mathcal{L} \\
& \left.\begin{array}{l}
\boldsymbol{A}^\top \boldsymbol{\alpha}(\boldsymbol{\ell}) = \displaystyle\sum_{k \in \mathcal{K}} \boldsymbol{\eta}^k(\boldsymbol{\ell}) \\
\boldsymbol{A}^\top \boldsymbol{\alpha}^k(\boldsymbol{\ell}) + [\boldsymbol{H}]_{\boldsymbol{\ell}_k(\boldsymbol{\ell})} \gamma_k(\boldsymbol{\ell}) + \boldsymbol{\eta}^k(\boldsymbol{\ell}) = \mathbf{0} \quad \forall k \in \mathcal{K} \\
\boldsymbol{b}^\top \left(\boldsymbol{\alpha}(\boldsymbol{\ell}) + \displaystyle\sum_{k \in \mathcal{K}} \boldsymbol{\alpha}^k(\boldsymbol{\ell}) \right) + \displaystyle\sum_{k \in \mathcal{K}} \left(\left[\boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^k \right]_{\boldsymbol{\ell}_k} - \epsilon \right) \gamma_k(\boldsymbol{\ell}) \leq -1
\end{array}\right\} \quad \forall \boldsymbol{\ell} \in \mathcal{L}_+.
\end{array}
\tag{EC.10}
$$

Since $\boldsymbol{\eta}^k(\boldsymbol{\ell})$ is free for all $k \in \mathcal{K}$ and $\boldsymbol{\ell} \in \mathcal{L}$, the first set of constraints associated with $\boldsymbol{\ell} \in \partial\mathcal{L}$ in (EC.10) is equivalent to

$$
\boldsymbol{A}^\top \left(\boldsymbol{\alpha}(\boldsymbol{\ell}) + \sum_{k \in \mathcal{K}} \boldsymbol{\alpha}^k(\boldsymbol{\ell}) \right) - \sum_{\substack{k \in \mathcal{K}: \\ \boldsymbol{\ell}_k = 0}} \boldsymbol{H}^\top \boldsymbol{\beta}^k(\boldsymbol{\ell}) + \sum_{\substack{k \in \mathcal{K}: \\ \boldsymbol{\ell}_k \neq 0}} [\boldsymbol{H}]_{\boldsymbol{\ell}_k} \gamma_k(\boldsymbol{\ell}) = \boldsymbol{C}\boldsymbol{x} + \sum_{k \in \mathcal{K}} \boldsymbol{\lambda}_k(\boldsymbol{\ell}) \cdot \boldsymbol{Q}\boldsymbol{y}^k,
$$

where we have exploited the fact that $\boldsymbol{\lambda}(\boldsymbol{\ell}) \geq \mathbf{0}$ and $\mathbf{e}^\top \boldsymbol{\lambda}(\boldsymbol{\ell}) = 1$. Similarly, the first set of constraints associated with $\boldsymbol{\ell} \in \mathcal{L}_+$ in (EC.10) is equivalent to

$$
\boldsymbol{A}^\top \left(\boldsymbol{\alpha}(\boldsymbol{\ell}) + \sum_{k \in \mathcal{K}} \boldsymbol{\alpha}^k(\boldsymbol{\ell}) \right) + \sum_{k \in \mathcal{K}} [\boldsymbol{H}]_{\boldsymbol{\ell}_k} \gamma_k(\boldsymbol{\ell}) = \mathbf{0}.
$$

We conclude that, in the presence of only exogenous uncertainty, Problem (13) reduces to

$$
\begin{array}{ll}
\min & \tau \\
\text{s.t.} & \tau \in \mathbb{R},\, \boldsymbol{x} \in \mathcal{X},\, \boldsymbol{y}^k \in \mathcal{Y},\, k \in \mathcal{K} \\
& \boldsymbol{\alpha}(\boldsymbol{\ell}) \in \mathbb{R}^R_+,\, \boldsymbol{\gamma}(\boldsymbol{\ell}) \in \mathbb{R}^K_+,\, \boldsymbol{\ell} \in \mathcal{L} \\
& \left.\begin{array}{l}
\boldsymbol{\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}) - \displaystyle\sum_{\substack{k \in \mathcal{K}: \\ \boldsymbol{\ell}_k = 0}} \boldsymbol{H}^\top \boldsymbol{\beta}^k(\boldsymbol{\ell}) + \sum_{\substack{k \in \mathcal{K}: \\ \boldsymbol{\ell}_k \neq 0}} [\boldsymbol{H}]_{\boldsymbol{\ell}_k} \gamma_k(\boldsymbol{\ell}) = \boldsymbol{C}\boldsymbol{x} + \sum_{k \in \mathcal{K}} \boldsymbol{\lambda}_k(\boldsymbol{\ell}) \cdot \boldsymbol{Q}\boldsymbol{y}^k \\
\tau \geq \boldsymbol{b}^\top \boldsymbol{\alpha}(\boldsymbol{\ell}) - \displaystyle\sum_{\substack{k \in \mathcal{K}: \\ \boldsymbol{\ell}_k = 0}} (\boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^k)^\top \boldsymbol{\beta}^k(\boldsymbol{\ell}) + \sum_{\substack{k \in \mathcal{K}: \\ \boldsymbol{\ell}_k \neq 0}} \left(\left[\boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^k\right]_{\boldsymbol{\ell}_k} - \epsilon \right) \gamma_k(\boldsymbol{\ell})
\end{array}\right\} \quad \forall \boldsymbol{\ell} \in \partial\mathcal{L} \\
& \left.\begin{array}{l}
\boldsymbol{A}^\top \boldsymbol{\alpha}(\boldsymbol{\ell}) + \displaystyle\sum_{k \in \mathcal{K}} [\boldsymbol{H}]_{\boldsymbol{\ell}_k} \gamma_k(\boldsymbol{\ell}) = \boldsymbol{0} \\
\boldsymbol{b}^\top \boldsymbol{\alpha}(\boldsymbol{\ell}) + \displaystyle\sum_{k \in \mathcal{K}} \left(\left[\boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^k\right]_{\boldsymbol{\ell}_k} - \epsilon \right) \gamma_k(\boldsymbol{\ell}) \leq -1
\end{array}\right\} \quad \forall \boldsymbol{\ell} \in \mathcal{L}_+,
\end{array}
\tag{EC.11}
$$

where we use the change of variables $\boldsymbol{\alpha}(\boldsymbol{\ell}) \leftarrow \left(\boldsymbol{\alpha}(\boldsymbol{\ell}) + \sum_{k \in \mathcal{K}} \boldsymbol{\alpha}^k(\boldsymbol{\ell})\right)$. We thus recover the MBLP formulation of the K-adaptability problem (6) from Hanasusanto et al. (2015), which concludes the proof. □

EC.5. Proofs of Statements in Section 6

Proof of Observation 3 Suppose that $\mathcal{X} := \{\boldsymbol{x} : \mathbf{e}^\top \boldsymbol{x} = 1\}$, $\mathcal{W} := \{\boldsymbol{w} : \mathbf{e}^\top \boldsymbol{w} = 1\}$, and $\mathcal{Y} := \{\boldsymbol{y} : \mathbf{e}^\top \boldsymbol{y} = 1\}$. Then,

$$
\begin{array}{lll}
\max & \boldsymbol{\xi}^\top \boldsymbol{C}\, \boldsymbol{x}' + \boldsymbol{\xi}^\top \boldsymbol{D}\, \boldsymbol{w}' + \boldsymbol{\xi}^\top \boldsymbol{Q}\, \boldsymbol{y}' & = \quad \displaystyle\max_{i,j,k} \; \left\{ \boldsymbol{\xi}^\top \boldsymbol{C}\, \mathbf{e}_i + \boldsymbol{\xi}^\top \boldsymbol{D}\, \mathbf{e}_j + \boldsymbol{\xi}^\top \boldsymbol{Q}\, \mathbf{e}_k \right\}. \\
\text{s.t.} & \boldsymbol{x}' \in \mathcal{X},\, \boldsymbol{w}' \in \mathcal{W},\, \boldsymbol{y}' \in \mathcal{Y} &
\end{array}
$$

Thus, in this case, the objective function is expressible in the form (14) and the claim follows. An analogous argument can be made if $\boldsymbol{C} = \boldsymbol{0}$, $\boldsymbol{D} = \boldsymbol{0}$, or $\boldsymbol{Q} = \boldsymbol{0}$. □

Proof of Lemma 3 It suffices to show that, for any fixed $\boldsymbol{x} \in \mathcal{X}$, $\boldsymbol{w} \in \mathcal{W}$, $\boldsymbol{y}^k \in \mathcal{Y}$, $k \in \mathcal{K}$, and $\overline{\boldsymbol{\xi}} \in \Xi$,

$$
\min_{k \in \mathcal{K}} \; \max_{\boldsymbol{\xi} \in \Xi(\boldsymbol{w}, \overline{\boldsymbol{\xi}})} \; \left\{ \max_{i \in \mathcal{I}} \; \boldsymbol{\xi}^\top \boldsymbol{C}^i\, \boldsymbol{x} + \boldsymbol{\xi}^\top \boldsymbol{D}^i\, \boldsymbol{w} + \boldsymbol{\xi}^\top \boldsymbol{Q}^i\, \boldsymbol{y}^k \right\} \tag{EC.12}
$$

and

$$
\max_{\substack{\boldsymbol{\xi}^k \in \Xi(\boldsymbol{w}, \overline{\boldsymbol{\xi}}), \\ k \in \mathcal{K}}} \; \min_{k \in \mathcal{K}} \; \left\{ \max_{i \in \mathcal{I}} \; (\boldsymbol{\xi}^k)^\top \boldsymbol{C}^i\, \boldsymbol{x} + (\boldsymbol{\xi}^k)^\top \boldsymbol{D}^i\, \boldsymbol{w} + (\boldsymbol{\xi}^k)^\top \boldsymbol{Q}^i\, \boldsymbol{y}^k \right\} \tag{EC.13}
$$

are equivalent.

First, note that Problem (EC.12) is always feasible and has a finite objective by virtue of the compactness of $\Xi(\boldsymbol{w}, \overline{\boldsymbol{\xi}})$ which is non-empty. Similarly, Problem (EC.13) is always feasible and has a finite objective.

We now show that both problems have the same objective. Let $\tilde{k}$ and $\{\tilde{\boldsymbol{\xi}}^k\}_{k\in\mathcal{K}}$ be feasible in (EC.12) and (EC.13), respectively. The objective value attained by $\tilde{k}$ in Problem (EC.12) is

$$\max_{\boldsymbol{\xi}\in\Xi(\boldsymbol{w},\overline{\boldsymbol{\xi}})} \left\{ \max_{i\in\mathcal{I}} \ \boldsymbol{\xi}^\top \boldsymbol{C}^i\,\boldsymbol{x} + \boldsymbol{\xi}^\top \boldsymbol{D}^i\,\boldsymbol{w} + \boldsymbol{\xi}^\top \boldsymbol{Q}^i\,\boldsymbol{y}^{\tilde{k}} \right\}.$$

Accordingly, the objective value attained by $\{\tilde{\boldsymbol{\xi}}^k\}_{k\in\mathcal{K}}$ in Problem (EC.13) is

$$\min_{k\in\mathcal{K}} \left\{ \max_{i\in\mathcal{I}} \ (\tilde{\boldsymbol{\xi}}^k)^\top \boldsymbol{C}^i\,\boldsymbol{x} + (\tilde{\boldsymbol{\xi}}^k)^\top \boldsymbol{D}^i\,\boldsymbol{w} + (\tilde{\boldsymbol{\xi}}^k)^\top \boldsymbol{Q}^i\,\boldsymbol{y}^k \right\}.$$

Note that

$$\begin{aligned} & \min_{k\in\mathcal{K}} \left\{ \max_{i\in\mathcal{I}} \ (\tilde{\boldsymbol{\xi}}^k)^\top \boldsymbol{C}^i\,\boldsymbol{x} + (\tilde{\boldsymbol{\xi}}^k)^\top \boldsymbol{D}^i\,\boldsymbol{w} + (\tilde{\boldsymbol{\xi}}^k)^\top \boldsymbol{Q}^i\,\boldsymbol{y}^k \right\} \\ \leq \quad & \max_{i\in\mathcal{I}} \ (\tilde{\boldsymbol{\xi}}^{\tilde{k}})^\top \boldsymbol{C}^i\,\boldsymbol{x} + (\tilde{\boldsymbol{\xi}}^{\tilde{k}})^\top \boldsymbol{D}^i\,\boldsymbol{w} + (\tilde{\boldsymbol{\xi}}^{\tilde{k}})^\top \boldsymbol{Q}^i\,\boldsymbol{y}^{\tilde{k}} \\ \leq \quad & \max_{\boldsymbol{\xi}\in\Xi(\boldsymbol{w},\overline{\boldsymbol{\xi}})} \ \max_{i\in\mathcal{I}} \ \boldsymbol{\xi}^\top \boldsymbol{C}\,\boldsymbol{x} + \boldsymbol{\xi}^\top \boldsymbol{D}\,\boldsymbol{w} + \boldsymbol{\xi}^\top \boldsymbol{Q}\,\boldsymbol{y}^{\tilde{k}}. \end{aligned}$$

Since the choice of $\tilde{k}\in\mathcal{K}$ and $\tilde{\boldsymbol{\xi}}^k \in \Xi(\boldsymbol{w},\overline{\boldsymbol{\xi}})$ was arbitrary, it follows that Problem (EC.12) upper bounds Problem (EC.13).

Next, we show that the converse also holds. Let

$$\boldsymbol{\xi}^{k,*} \in \arg\max_{\boldsymbol{\xi}\in\Xi(\boldsymbol{w},\overline{\boldsymbol{\xi}})} \left\{ \max_{i\in\mathcal{I}} \ \boldsymbol{\xi}^\top \boldsymbol{C}^i\,\boldsymbol{x} + \boldsymbol{\xi}^\top \boldsymbol{D}^i\,\boldsymbol{w} + \boldsymbol{\xi}^\top \boldsymbol{Q}^i\,\boldsymbol{y}^k \right\}.$$

Then, the optimal objective value of Problem (EC.12) is expressible as

$$\min_{k\in\mathcal{K}} \left\{ \max_{i\in\mathcal{I}} \ (\boldsymbol{\xi}^{k,*})^\top \boldsymbol{C}^i\,\boldsymbol{x} + (\boldsymbol{\xi}^{k,*})^\top \boldsymbol{D}^i\,\boldsymbol{w} + (\boldsymbol{\xi}^{k,*})^\top \boldsymbol{Q}^i\,\boldsymbol{y}^k \right\}.$$

The solution $\{\boldsymbol{\xi}^{k,*}\}_{k\in\mathcal{K}}$ is feasible in (EC.13) with objective

$$\min_{k\in\mathcal{K}} \left\{ \max_{i\in\mathcal{I}} \ (\boldsymbol{\xi}^{k,*})^\top \boldsymbol{C}^i\,\boldsymbol{x} + (\boldsymbol{\xi}^{k,*})^\top \boldsymbol{D}^i\,\boldsymbol{w} + (\boldsymbol{\xi}^{k,*})^\top \boldsymbol{Q}^i\,\boldsymbol{y}^k \right\}.$$

Thus, the optimal objective value of Problem (EC.13) upper bounds that of Problem (EC.12).

Combining the two parts of the proof, we conclude that Problems (EC.12) and (EC.13) are equivalent. □

Proof of Theorem 6 The objective function of Problem (15) is expressible as

$$\max_{\overline{\boldsymbol{\xi}}\in\Xi} \ \max_{\boldsymbol{\xi}^k\in\Xi(\boldsymbol{w},\overline{\boldsymbol{\xi}}),k\in\mathcal{K}} \ \min_{k\in\mathcal{K}} \left\{ \max_{i\in\mathcal{I}} \ \boldsymbol{\xi}^\top \boldsymbol{C}^i\,\boldsymbol{x} + \boldsymbol{\xi}^\top \boldsymbol{D}^i\,\boldsymbol{w} + \boldsymbol{\xi}^\top \boldsymbol{Q}^i\,\boldsymbol{y}^k \right\}.$$

Using an epigraph reformulation, we can write it equivalently as

$$\begin{array}{ll} \text{maximize} & \tau \\ \text{subject to} & \tau\in\mathbb{R},\ \overline{\boldsymbol{\xi}}\in\Xi,\ \boldsymbol{\xi}^k\in\Xi(\boldsymbol{w},\overline{\boldsymbol{\xi}}),\ k\in\mathcal{K} \\ & \tau \leq \max_{i\in\mathcal{I}} \ (\boldsymbol{\xi}^k)^\top \boldsymbol{C}^i\,\boldsymbol{x} + (\boldsymbol{\xi}^k)^\top \boldsymbol{D}^i\,\boldsymbol{w} + (\boldsymbol{\xi}^k)^\top \boldsymbol{Q}^i\,\boldsymbol{y}^k \quad \forall k\in\mathcal{K}. \end{array} \tag{EC.14}$$

Noting that, for each $k \in \mathcal{K}$, the choice of $i \in \mathcal{K}$ can be made, in conjuction with the choice in τ, $\overline{\boldsymbol{\xi}}$, and $\boldsymbol{\xi}^k$, $k \in \mathcal{K}$, Problem (EC.14) can be written equivalently as

$$\begin{array}{lll} \underset{\boldsymbol{i}_k \in \mathcal{I},\, k\in\mathcal{K}}{\text{maximize}} & \max & \tau \\ & \text{s.t.} & \tau \in \mathbb{R},\, \overline{\boldsymbol{\xi}} \in \Xi,\, \boldsymbol{\xi}^k \in \Xi(\boldsymbol{w}, \overline{\boldsymbol{\xi}}),\, \forall k \in \mathcal{K} \\ & & \tau \,\leq\, (\boldsymbol{\xi}^k)^\top \boldsymbol{C}^{\boldsymbol{i}_k} \boldsymbol{x} + (\boldsymbol{\xi}^k)^\top \boldsymbol{D}^{\boldsymbol{i}_k} \boldsymbol{w} + (\boldsymbol{\xi}^k)^\top \boldsymbol{Q}^{\boldsymbol{i}_k} \boldsymbol{y}^k \quad \forall k \in \mathcal{K}. \end{array} \tag{EC.15}$$

Dualizing the inner maximization problem yields

$$\begin{array}{lll} \underset{\boldsymbol{i}_k \in \mathcal{I},\, k\in\mathcal{K}}{\text{maximize}} & \min & \boldsymbol{b}^\top \boldsymbol{\beta} + \sum_{k\in\mathcal{K}} \boldsymbol{b}^\top \boldsymbol{\beta}^k \\ & \text{s.t.} & \boldsymbol{\alpha} \in \mathbb{R}^K_+,\, \boldsymbol{\beta} \in \mathbb{R}^R_+,\, \boldsymbol{\beta}^k \in \mathbb{R}^R_+,\, \boldsymbol{\gamma}^k \in \mathbb{R}^{N_\xi},\, \forall k \in \mathcal{K} \\ & & \mathbf{e}^\top \boldsymbol{\alpha} = 1 \\ & & \boldsymbol{A}^\top \boldsymbol{\beta}^k + \boldsymbol{w} \circ \boldsymbol{\gamma}^k = \boldsymbol{\alpha}_k \left(\boldsymbol{C}^{\boldsymbol{i}_k} \boldsymbol{x} + \boldsymbol{D}^{\boldsymbol{i}_k} \boldsymbol{w} + \boldsymbol{Q}^{\boldsymbol{i}_k} \boldsymbol{y}^k \right) \quad \forall k \in \mathcal{K} \\ & & \boldsymbol{A}^\top \boldsymbol{\beta} = \sum_{k\in\mathcal{K}} \boldsymbol{w} \circ \boldsymbol{\gamma}^k. \end{array} \tag{EC.16}$$

Equivalence of Problems (EC.15) and (EC.16) follows by strong LP duality which applies since the inner maximization problem in (EC.15) is feasible and bounded. We next interchange the max and min operators, indexing each of the decision variables by $\boldsymbol{i} := (\boldsymbol{i}_1, \ldots, \boldsymbol{i}_k) \in \mathcal{I}^K$. We obtain

$$\begin{array}{ll} \text{minimize} & \max\limits_{\boldsymbol{i}\in\mathcal{I}^K} \;\; \boldsymbol{b}^\top \boldsymbol{\beta}^{\boldsymbol{i}} + \sum\limits_{k\in\mathcal{K}} \boldsymbol{b}^\top \boldsymbol{\beta}^{\boldsymbol{i},k} \\ \text{subject to} & \boldsymbol{\alpha}^{\boldsymbol{i}} \in \mathbb{R}^K_+,\, \boldsymbol{\beta}^{\boldsymbol{i}} \in \mathbb{R}^R_+,\, \boldsymbol{\beta}^{\boldsymbol{i},k} \in \mathbb{R}^R_+,\, \boldsymbol{\gamma}^{\boldsymbol{i},k} \in \mathbb{R}^{N_\xi},\, \forall k \in \mathcal{K},\, \boldsymbol{i} \in \mathcal{I}^K \\ & \left. \begin{array}{l} \mathbf{e}^\top \boldsymbol{\alpha}^{\boldsymbol{i}} = 1 \\ \boldsymbol{A}^\top \boldsymbol{\beta}^{\boldsymbol{i},k} + \boldsymbol{w} \circ \boldsymbol{\gamma}^{\boldsymbol{i},k} = \boldsymbol{\alpha}^{\boldsymbol{i}}_k \left(\boldsymbol{C}^{\boldsymbol{i}_k} \boldsymbol{x} + \boldsymbol{D}^{\boldsymbol{i}_k} \boldsymbol{w} + \boldsymbol{Q}^{\boldsymbol{i}_k} \boldsymbol{y}^k \right) \quad \forall k \in \mathcal{K} \\ \boldsymbol{A}^\top \boldsymbol{\beta}^{\boldsymbol{i}} = \sum\limits_{k\in\mathcal{K}} \boldsymbol{w} \circ \boldsymbol{\gamma}^{\boldsymbol{i},k} \end{array} \right\} \quad \forall \boldsymbol{i} \in \mathcal{I}^K. \end{array}$$

Finally, we write the above problem as a single minimization using an epigraph formulation, as follows

$$\begin{array}{ll} \text{minimize} & \tau \\ \text{subject to} & \tau \in \mathbb{R},\, \boldsymbol{\alpha}^{\boldsymbol{i}} \in \mathbb{R}^K_+,\, \boldsymbol{\beta}^{\boldsymbol{i}} \in \mathbb{R}^R_+,\, \boldsymbol{\beta}^{\boldsymbol{i},k} \in \mathbb{R}^R_+,\, \boldsymbol{\gamma}^{\boldsymbol{i},k} \in \mathbb{R}^{N_\xi},\, \forall k \in \mathcal{K},\, \boldsymbol{i} \in \mathcal{I}^K \\ & \left. \begin{array}{l} \tau \,\geq\, \boldsymbol{b}^\top \boldsymbol{\beta}^{\boldsymbol{i}} + \sum_{k\in\mathcal{K}} \boldsymbol{b}^\top \boldsymbol{\beta}^{\boldsymbol{i},k} \\ \mathbf{e}^\top \boldsymbol{\alpha}^{\boldsymbol{i}} = 1 \\ \boldsymbol{A}^\top \boldsymbol{\beta}^{\boldsymbol{i},k} + \boldsymbol{w} \circ \boldsymbol{\gamma}^{\boldsymbol{i},k} = \boldsymbol{\alpha}^{\boldsymbol{i}}_k \left(\boldsymbol{C}^{\boldsymbol{i}_k} \boldsymbol{x} + \boldsymbol{D}^{\boldsymbol{i}_k} \boldsymbol{w} + \boldsymbol{Q}^{\boldsymbol{i}_k} \boldsymbol{y}^k \right) \quad \forall k \in \mathcal{K} \\ \boldsymbol{A}^\top \boldsymbol{\beta}^{\boldsymbol{i}} = \sum\limits_{k\in\mathcal{K}} \boldsymbol{w} \circ \boldsymbol{\gamma}^{\boldsymbol{i},k} \end{array} \right\} \quad \forall \boldsymbol{i} \in \mathcal{I}^K. \end{array} \tag{EC.17}$$

The claim then follows by grouping the outer minimization problem in (15) with the minimization problem in (EC.17). $\square$

Proof of Proposition 2 Since $(\boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}^k\}_{k\in\mathcal{K}})$ is feasible in the relaxed master problem $(\mathcal{CCG}_{\mathrm{mstr}}(\widetilde{\mathcal{I}}))$, it follows that $\boldsymbol{x} \in \mathcal{X}$, $\boldsymbol{w} \in \mathcal{W}$, and $\boldsymbol{y}^k \in \mathcal{Y}$, $k \in \mathcal{K}$. Thus, $(\boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}^k\}_{k\in\mathcal{K}})$ is feasible in Problem $(\mathcal{PO}_K^{\mathrm{PWL}})$. An inspection of the Proof of Theorem 6 reveals that the objective value of $(\boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}^k\}_{k\in\mathcal{K}})$ in Problem $(\mathcal{PO}_K^{\mathrm{PWL}})$ is given by the optimal value of Problem (EC.14). The proof then follows by noting that Problems (EC.14) and $(\mathcal{CCG}_{\mathrm{feas}}(\boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}^k\}_{k\in\mathcal{K}}))$ are equivalent. $\square$

Proof of Lemma 4 *(i)* By virtue of Proposition 2, it follows that $\theta \geq \tau$.

(ii) Suppose that $\theta = \tau$ and that there exists $\boldsymbol{i} \in \mathcal{I}^K$ such that Problem $(\mathcal{CCG}^{\boldsymbol{i}}_{\mathrm{sub}}(\tau, \boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}^k\}_{k\in\mathcal{K}}))$ is infeasible. This implies that there exists $\boldsymbol{i} \in \mathcal{I}^K$ such that τ is strictly smaller than the optimal objective value of

$$
\begin{array}{ll}
\text{minimize} & \boldsymbol{b}^\top \boldsymbol{\beta}^{\boldsymbol{i}} + \sum_{k\in\mathcal{K}} \boldsymbol{b}^\top \boldsymbol{\beta}^{\boldsymbol{i},k} \\
\text{subject to} & \boldsymbol{\alpha}^{\boldsymbol{i}} \in \mathbb{R}^K_+,\ \boldsymbol{\beta}^{\boldsymbol{i}} \in \mathbb{R}^R_+,\ \boldsymbol{\beta}^{\boldsymbol{i},k} \in \mathbb{R}^R_+,\ \boldsymbol{\gamma}^{\boldsymbol{i},k} \in \mathbb{R}^{N_\xi},\ \forall k \in \mathcal{K} \\
& \mathbf{e}^\top \boldsymbol{\alpha}^{\boldsymbol{i}} = 1 \\
& \boldsymbol{A}^\top \boldsymbol{\beta}^{\boldsymbol{i},k} + \boldsymbol{w} \circ \boldsymbol{\gamma}^{\boldsymbol{i},k} = \boldsymbol{\alpha}^{\boldsymbol{i}}_k \left(\boldsymbol{C}^{\boldsymbol{i}_k} \boldsymbol{x} + \boldsymbol{D}^{\boldsymbol{i}_k} \boldsymbol{w} + \boldsymbol{Q}^{\boldsymbol{i}_k} \boldsymbol{y}^k \right) \quad \forall k \in \mathcal{K} \\
& \boldsymbol{A}^\top \boldsymbol{\beta}^{\boldsymbol{i}} = \displaystyle\sum_{k\in\mathcal{K}} \boldsymbol{w} \circ \boldsymbol{\gamma}^{\boldsymbol{i},k}.
\end{array}
\tag{EC.18}
$$

Equivalently, by dualizing this problem, we conclude that there exists $\boldsymbol{i} \in \mathcal{I}^K$ such that τ is strictly smaller than the optimal objective value of

$$
\begin{array}{ll}
\text{maximize} & \theta' \\
\text{subject to} & \theta' \in \mathbb{R},\ \overline{\boldsymbol{\xi}} \in \Xi,\ \boldsymbol{\xi}^k \in \Xi(\boldsymbol{w}, \overline{\boldsymbol{\xi}}),\ \forall k \in \mathcal{K} \\
& \theta' \leq (\boldsymbol{\xi}^k)^\top \boldsymbol{C}^{\boldsymbol{i}_k} \boldsymbol{x} + (\boldsymbol{\xi}^k)^\top \boldsymbol{D}^{\boldsymbol{i}_k} \boldsymbol{w} + (\boldsymbol{\xi}^k)^\top \boldsymbol{Q}^{\boldsymbol{i}_k} \boldsymbol{y}^k \quad \forall k \in \mathcal{K}.
\end{array}
\tag{EC.19}
$$

Since Problem (EC.19) lower bounds Problem $(\mathcal{CCG}_{\mathrm{feas}}(\boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}^k\}_{k\in\mathcal{K}}))$ with optimal objective value θ, we conclude that $\tau < \theta$, a contradiction.

(iii) Suppose that $\theta > \tau$ and let $\boldsymbol{i}$ be defined as in the premise of the lemma. Then, $\boldsymbol{i}$ is optimal in (EC.16) with associated optimal objective value θ. This implies that the optimal objective value of Problem (EC.18) is θ. Since $\theta > \tau$, this implies that subproblem $(\mathcal{CCG}^{\boldsymbol{i}}_{\mathrm{sub}}(\tau, \boldsymbol{x}, \boldsymbol{w}, \{\boldsymbol{y}^k\}_{k\in\mathcal{K}}))$ is infeasible, which concludes the proof.

We have thus proved all claims. $\square$

Proof of Theorem 7 First, note that finite termination is guaranteed since at each iteration, either $\mathrm{UB} - \mathrm{LB} \leq \delta$ (in which case the algorithm terminates) or a new set of constraints (indexed by the infeasible index $\boldsymbol{i}$) is added to the master problem ($\mathcal{CCG}_{\rm mstr}(\widetilde{\mathcal{I}})$), see Lemma 4. Since the set of all indices, $\mathcal{I}^K$, is finite, the algorithm will terminate in a finite number of steps. Second, by construction, at any iteration of the algorithm, τ (i.e., LB) provides a lower bound on the optimal objective value of the problem. On the other hand, the returned (feasible) solution has as objective value θ (i.e., UB). Since the algorithm only terminates if $\text{UB} - \text{LB} \leq \delta$, we are guaranteed that, at termination, the returned solution will have an objective value that is within δ of the optimal objective value of the problem. This concludes the proof. $\square$

EC.6. Proofs of Statements in Section 7

Proof of Theorem 10 For any fixed $\{\boldsymbol{w}^{t,k_t}\}_{t\in\mathcal{T},k_t\in\mathcal{K}}$ and $\{\boldsymbol{y}^{t,k_t}\}_{t\in\mathcal{T},k_t\in\mathcal{K}}$, the inner problem in the objective of Problem (22) can be written in epigraph form as

$$
\begin{array}{ll}
\text{maximize} & \tau \\
\text{subject to} & \tau \in \mathbb{R},\ \boldsymbol{\xi}^{T,k_T\cdots k_1} \in \Xi^T(\boldsymbol{w}^{1,k_1},\ldots,\boldsymbol{w}^{T-1,k_{T-1}}) \quad \forall k_1,\ldots,k_T \in \mathcal{K} \\
& \tau \leq \displaystyle\sum_{t\in\mathcal{T}} (\boldsymbol{\xi}^{T,k_T\cdots k_1})^\top \boldsymbol{D}^t\, \boldsymbol{w}^{t,k_t} + (\boldsymbol{\xi}^{T,k_T\cdots k_1})^\top \boldsymbol{Q}^t\, \boldsymbol{y}^{t,k_t} \quad \forall k_1,\ldots,k_T \in \mathcal{K}.
\end{array}
$$

From the definition of $\Xi^T(\cdot)$ in Lemma 5, the above problem can be equivalently written as

$$
\begin{array}{ll}
\text{maximize} & \tau \\
\text{subject to} & \tau \in \mathbb{R},\ \boldsymbol{\xi}^{t,k_t\cdots k_1} \in \Xi \quad \forall t \in \mathcal{T},\ k_1,\ldots,k_t \in \mathcal{K} \\
& \tau \leq \displaystyle\sum_{t\in\mathcal{T}} (\boldsymbol{\xi}^{T,k_T\cdots k_1})^\top \boldsymbol{D}^t\, \boldsymbol{w}^{t,k_t} + (\boldsymbol{\xi}^{T,k_T\cdots k_1})^\top \boldsymbol{Q}^t\, \boldsymbol{y}^{t,k_t} \quad \forall k_1,\ldots,k_T \in \mathcal{K} \\
& \boldsymbol{w}^{t-1,k_{t-1}} \circ \boldsymbol{\xi}^{t,k_t\cdots k_1} = \boldsymbol{w}^{t-1,k_{t-1}} \circ \boldsymbol{\xi}^{t-1,k_{t-1}\cdots k_1} \quad \forall t \in \mathcal{T}\backslash\{1\},\ k_1,\ldots,k_t \in \mathcal{K}.
\end{array}
$$

Writing the set Ξ explicitly yields

$$
\begin{array}{ll}
\text{maximize} & \tau \\
\text{subject to} & \tau \in \mathbb{R},\ \boldsymbol{\xi}^{t,k_t\cdots k_1} \in \mathbb{R}^{N_\xi} \quad \forall t \in \mathcal{T},\ k_1,\ldots,k_t \in \mathcal{K} \\
& \tau \leq \displaystyle\sum_{t\in\mathcal{T}} \left(\boldsymbol{D}^t\, \boldsymbol{w}^{t,k_t} + \boldsymbol{Q}^t\, \boldsymbol{y}^{t,k_t}\right)^\top \boldsymbol{\xi}^{T,k_T\cdots k_1} \quad \forall k_1,\ldots,k_T \in \mathcal{K} \\
& \boldsymbol{A}\boldsymbol{\xi}^{t,k_t\cdots k_1} \leq \boldsymbol{b} \quad \forall t \in \mathcal{T},\ k_1,\ldots,k_t \in \mathcal{K} \\
& \boldsymbol{w}^{t-1,k_{t-1}} \circ \boldsymbol{\xi}^{t,k_t\cdots k_1} = \boldsymbol{w}^{t-1,k_{t-1}} \circ \boldsymbol{\xi}^{t-1,k_{t-1}\cdots k_1} \quad \forall t \in \mathcal{T}\backslash\{1\},\ k_1,\ldots,k_t \in \mathcal{K}.
\end{array}
$$

The dual of this problem reads

$$
\begin{array}{ll}
\text{minimize} & \displaystyle\sum_{t\in\mathcal{T}}\sum_{k_1\in\mathcal{K}}\cdots\sum_{k_t\in\mathcal{K}} \boldsymbol{b}^\top\boldsymbol{\beta}^{t,k_1\cdots k_t} \\
\text{subject to} & \boldsymbol{\alpha}\in\mathbb{R}_+^{K^T},\ \boldsymbol{\beta}^{t,k_1\cdots k_t}\in\mathbb{R}_+^R,\ \boldsymbol{\gamma}^{t,k_1\cdots k_t}\in\mathbb{R}^{N_\xi},\ t\in\mathcal{T},\ k_1,\ldots,k_t\in\mathcal{K} \\
& \mathbf{e}^\top\boldsymbol{\alpha}=1 \\
& \boldsymbol{A}^\top\boldsymbol{\beta}^{1,k_1}=\displaystyle\sum_{k_2\in\mathcal{K}}\boldsymbol{w}^{1,k_1}\circ\boldsymbol{\gamma}^{2,k_1k_2}\quad\forall k_1\in\mathcal{K} \\
& \boldsymbol{A}^\top\boldsymbol{\beta}^{t,k_1\cdots k_t}+\boldsymbol{w}^{t-1,k_{t-1}}\circ\boldsymbol{\gamma}^{t,k_1\cdots k_t}=\displaystyle\sum_{k_{t+1}\in\mathcal{K}}\boldsymbol{w}^{t,k_t}\circ\boldsymbol{\gamma}^{t+1,k_1\cdots k_{t+1}}\quad\forall t\in\mathcal{T}\backslash\{1,T\},\ k_1,\ldots,k_t\in\mathcal{K} \\
& \boldsymbol{A}^\top\boldsymbol{\beta}^{T,k_1\cdots k_T}+\boldsymbol{w}^{T-1,k_{T-1}}\circ\boldsymbol{\gamma}^{T,k_1\cdots k_T}\ =\ \boldsymbol{\alpha}_{k_1\cdots k_T}\displaystyle\sum_{t\in\mathcal{T}}\left(\boldsymbol{D}^t\,\boldsymbol{w}^{t,k_t}+\boldsymbol{Q}^t\,\boldsymbol{y}^{t,k_t}\right)\quad\forall k_1,\ldots k_T.
\end{array}
$$

Moreover, strong duality applies by virtue of the compactness of Ξ. Merging the problem above with the outer minimization problem in (22) yields

$$
\begin{array}{ll}
\text{minimize} & \displaystyle\sum_{t\in\mathcal{T}}\sum_{k_1\in\mathcal{K}}\cdots\sum_{k_t\in\mathcal{K}} \boldsymbol{b}^\top\boldsymbol{\gamma}^{t,k_1\cdots k_t} \\
\text{subject to} & \boldsymbol{\alpha}\in\mathbb{R}_+^{K^T},\ \boldsymbol{\beta}^{t,k_1\cdots k_t}\in\mathbb{R}_+^R,\ \boldsymbol{\gamma}^{t,k_1\cdots k_t}\in\mathbb{R}^{N_\xi},\ t\in\mathcal{T},\ k_1,\ldots,k_t\in\mathcal{K} \\
& \boldsymbol{y}^{t,k}\in\mathcal{Y}_t,\ \boldsymbol{w}^{t,k}\in\mathcal{W}_t\quad\forall t\in\mathcal{T},\ k\in\mathcal{K} \\
& \mathbf{e}^\top\boldsymbol{\alpha}=1 \\
& \boldsymbol{A}^\top\boldsymbol{\beta}^{1,k_1}=\displaystyle\sum_{k_2\in\mathcal{K}}\boldsymbol{w}^{1,k_1}\circ\boldsymbol{\gamma}^{2,k_1k_2}\quad\forall k_1\in\mathcal{K} \\
& \boldsymbol{A}^\top\boldsymbol{\beta}^{t,k_1\cdots k_t}+\boldsymbol{w}^{t-1,k_{t-1}}\circ\boldsymbol{\gamma}^{t,k_1\cdots k_t}=\displaystyle\sum_{k_{t+1}\in\mathcal{K}}\boldsymbol{w}^{t,k_t}\circ\boldsymbol{\gamma}^{t+1,k_1\cdots k_{t+1}}\quad\forall t\in\mathcal{T}\backslash\{1,T\},\ k_1,\ldots,k_t\in\mathcal{K} \\
& \boldsymbol{A}^\top\boldsymbol{\beta}^{T,k_1\cdots k_T}+\boldsymbol{w}^{T-1,k_{T-1}}\circ\boldsymbol{\gamma}^{T,k_1\cdots k_T}\ =\ \boldsymbol{\alpha}_{k_1\cdots k_T}\displaystyle\sum_{t\in\mathcal{T}}\left(\boldsymbol{D}^t\,\boldsymbol{w}^{t,k_t}+\boldsymbol{Q}^t\,\boldsymbol{y}^{t,k_t}\right)\quad\forall k_1,\ldots k_T\in\mathcal{K} \\
& \boldsymbol{w}^{t,k_t}\geq\boldsymbol{w}^{t-1,k_{t-1}}\quad\forall t\in\mathcal{T},\ k_t\in\mathcal{K},\ k_{t-1}\in\mathcal{K} \\
& \displaystyle\sum_{t\in\mathcal{T}}\boldsymbol{V}^t\boldsymbol{w}^{t,k_t}+\boldsymbol{W}^t\boldsymbol{y}^{t,k_t}\leq\boldsymbol{h}\quad\forall k_1,\ldots,k_T\in\mathcal{K},
\end{array}
$$

and our proof is complete. $\square$

EC.7. Detailed Evaluation Results of Symmetry Breaking & Greedy Heuristic

The details of the evaluation results of symmetry breaking constraints and greedy heuristic approach on a synthetic dataset with $I=40$ items and $J=20$ features ($\Gamma=0$) are provided in Table EC.1. The table shows, for Q varied in the set $\{2,4,6,8\}$ and for $K=1,\ldots,10$: *a)* the objective value and solver time of this instance of Problem (11); *b)* the objective value and solver time of this instance of Problem (11) augmented with the symmetry breaking constraints presented in Section EC.1.1; *c)* the objective value and solver time of this instance of Problem (11) solved with the greedy heuristic approach presented in Section EC.1.2. Finally, the table summarizes, for each instance, the speed-up factor due to employing symmetry breaking and the loss in optimality due to employing the greedy heuristic. The speed-up factor is computed as the ratio of the solver time of the MILP without and with symmetry breaking constraints. The optimality gap of the greedy solution is computed as the gap between the objective value of the greedy solution and the objective value of Problem ($\mathcal{WCU}^{\text{PE}}$). A summary of these results is provided in Table 2.

Table EC.1 Detail of evaluation results of symmetry breaking constraints and greedy heuristic approach on a synthetic dataset with $I = 40$ items and $J = 20$ features ($\Gamma = 0$).

Q	K	No Symm. Break.		Symm. Break.		Heuristic		Statistics	
		Objective Value	Solve Time (sec)	Objective Value	Solve Time (sec)	Objective Value	Solve Time (sec)	Symm. Break. Speed-Up	Opt. Gap of Heuristic
2	1	0.196	0.492	0.196	0.258	0.196	0.216	1.907	0.000
2	2	0.348	2.105	0.348	0.785	0.307	0.428	2.682	11.782
2	3	0.416	48.994	0.416	9.999	0.395	1.225	4.900	5.048
2	4	0.416	5880.26	0.416	135.326	0.395	51.117	43.453	5.048
2	5	-	-	0.416	7144.08	0.395	65.333	-	5.048
2	6	-	-	-	-	0.395	91.773	-	-
2	7	-	-	-	-	0.395	147.397	-	-
2	8	-	-	-	-	0.395	134.338	-	-
2	9	-	-	-	-	0.395	165.555	-	-
2	10	-	-	-	-	0.395	226.524	-	-
4	1	0.196	0.595	0.196	0.584	0.196	0.181	1.019	0.000
4	2	0.348	1.603	0.348	1.528	0.307	0.378	1.049	11.782
4	3	0.416	31.63	0.416	13.311	0.395	1.332	2.376	5.048
4	4	0.435	1496.96	0.435	104.435	0.431	2.122	14.334	0.920
4	5	-	-	0.449	727.286	0.444	2.799	-	1.114
4	6	-	-	-	-	0.444	1147.550	-	-
4	7	-	-	-	-	0.444	2030.130	-	-
4	8	-	-	-	-	0.445	2139.650	-	-
4	9	-	-	-	-	0.445	4011.340	-	-
4	10	-	-	-	-	-	-	-	-
6	1	0.196	0.636	0.196	0.631	0.196	0.185	1.008	0.000
6	2	0.348	1.763	0.348	1.243	0.307	0.326	1.418	11.782
6	3	0.416	39.369	0.416	13.454	0.395	1.240	2.926	5.048
6	4	0.435	1210.09	0.435	102.138	0.431	1.759	11.848	0.920
6	5	-	-	0.449	883.301	0.444	3.277	-	1.114
6	6	-	-	-	-	0.450	9.400	-	-
6	7	-	-	-	-	0.463	15.919	-	-
6	8	-	-	-	-	0.471	4049.04	-	-
6	9	-	-	-	-	-	-	-	-
6	10	-	-	-	-	-	-	-	-
8	1	0.196	0.623	0.196	0.587	0.196	0.218	1.061	0.000
8	2	0.348	1.504	0.348	1.351	0.307	0.426	1.113	11.782
8	3	0.416	36.891	0.416	11.999	0.395	1.037	3.075	5.048
8	4	0.435	1913.76	0.435	78.795	0.431	1.879	24.288	0.920
8	5	-	-	0.449	702.598	0.444	5.912	-	1.114
8	6	-	-	0.465	3561.5	0.450	5.609	-	3.226
8	7	-	-	-	-	0.463	13.856	-	-
8	8	-	-	-	-	0.474	12.946	-	-
8	9	-	-	-	-	0.477	41.754	-	-
8	10	-	-	-	-	-	-	-	-