Active Preference Elicitation via Adjustable Robust Optimization

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

Duncan C. McElfresh
University of Maryland, Department of Mathematics, dmcelfre@math.umd.edu

Yingxiao Ye
University of Southern California, Center for Artificial Intelligence in Society, yeyingxi@usc.edu

John P. Dickerson
University of Maryland, Department of Computer Science, john@cs.umd.edu

Eric Rice
University of Southern California, Center for Artificial Intelligence in Society, erice@usc.edu

We consider the problem faced by a recommender system which seeks to offer a user with unknown preferences an item from a potentially uncountable or countably infinite 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 pairwise comparison, in the spirit of choice-based conjoint analysis. We take the point of view of either a risk averse or regret averse recommender system which only possess limited, set-based information on the user utility function. We investigate two settings: *a)* an *offline elicitation* setting, where all queries are made at once, and *b)* an *online elicitation* setting, where queries are selected sequentially over time in an adaptive fashion. We propose exact robust optimization formulations of these problems which integrate the preference elicitation and recommendation phases and study the complexity of these problems. For the offline case, where the active preference elicitation problem takes the form of a *two-stage* robust optimization problem with decision-dependent information discovery, we provide an enumeration-based algorithm and also an equivalent reformulation in the form of a mixed-binary linear program which we solve via column-and-constraint generation. For the online setting, where the active preference learning problem takes the form of a *multi-stage* robust optimization problem with decision-dependent information discovery, we propose a conservative solution approach. We evaluate the performance of our methods on both synthetic data and real data from the Homeless Management Information System. We simulate elicitation of the preferences of policy-makers in terms of characteristics of housing allocation policies (measures of fairness, efficiency, and interpretability) to better match individuals experiencing homelessness to scarce housing resources. Our framework is shown to outperform the state-of-the-art techniques from the literature.

Key words: robust optimization, decision-dependent information discovery, preference elicitation.

1. Introduction

1.1. Background

Automated decision-making systems (also known as *decision-support*, *recommender*, or *recommendation systems*) are increasingly being used to assist human decision-makers. In particular, they have become prevalent in environments where the set of alternatives is too large to enumerate (being potentially even uncountable or countably infinite), or when decisions need to be made in the presence of uncertainty or complex constraints. In such settings, several research studies have shown that human decision-makers often fail to identify an optimal (or near optimal) decision: they are *boundedly rational*, being limited by the information they have, their cognitive abilities, and the finite amount of time they have to make a decision, see Simon (1955). Automated decision-making systems on the other hand have proved extremely powerful at parsing through myriad alternatives in reasonable time to identify “the best” option. As an example, automated decision-making systems are routinely used to recommend routes (see e.g., Google[1] or Apple[2] maps).

To be able to identify the “best” alternative to recommend, the decision-support system needs to understand the *preferences* of the user (or *agent*) it is looking to assist. At the same time, preferences over alternatives can vary wildly from user to user. Thus, *preference modelling* and *preference elicitation* techniques are needed to be able to make *personalized recommendations* that users are likely to adopt. The problem of modeling and eliciting preferences has been of long standing interest in decision theory, psychology, and economics (see e.g., Neumann and Morgenstern (1944), Allais (1953), Debreu (1954), Slovic et al. (1977), and Kahneman and Tversky (1979)), operations research and management science (see e.g., Zionts and Wallenius (1976) and Dyer and Sarin (1979)), marketing (see e.g., Johnson (1987, 1991), Green and Srinivasan (1990), Carroll and Green (1997), and Toubia et al. (2003, 2004, 2007)), and more recently artificial intelligence (see e.g., Wang and Boutilier (2003), Boutilier et al. (2006) and Domshlak et al. (2011)).

The most common way to represent user preferences over a choice set is by means of the *binary preference relation*. Given two alternatives, this relation asks which one is preferred. This relation then induces a partial preorder over the elements of the choice set, allowing the automated decision-support system to *reason* over preferences, to be able to answer ordering style queries. For example, it can identify the preferred item from the choice set; or recommend a collection of preferred alternatives. Unfortunately, as the cardinality of the choice set grows, eliciting the preferences of agents over this set under the binary preference model becomes impracticable or extremely costly (this is sometimes referred to as the *preference bottleneck*) and a more structured framework is needed to achieve tractability.

Multi-attribute utility theory (see e.g., Dyer et al. (1992) and Keeney et al. (1993)) suggests that, when choosing one alternative over another, a user is basing their decision on the *attributes* of the two options. Under this model, alternatives are uniquely characterised by their attributes and can therefore be mapped to points in a (potentially high-dimensional) Cartesian space. Under certain mild assumptions, see Neumann and Morgenstern (1944) and Debreu (1954), the preferences of agents over goods under this model can be represented by a numeric function (not necessarily unique), termed *utility* or *value* function. This function maps the attributes of each item to a single number, such that an item is preferred over another item if and only if the former has a higher utility. This representation is useful to conduct preference elicitation since binary preference relations between options for example can be mapped back to relations between the attributes of the options, making the elicitation task more tractable.

Preference elicitation techniques (known as *conjoint analysis* in marketing) have emerged as a promising means to learn the decision-maker's preferences before making a (personalized) recommendation. They consist in *interactively querying* the user, through a moderate number of strategically chosen questions, each requiring little cognitive effort, to then be able to make high quality recommendations based on the preference estimates. The most common type of query is a *comparison query* (referred to as *choice-based conjoint* technique in the marketing literature). It takes the form of a pairwise comparison between alternatives, e.g., "Do you prefer option A or option B?". More sophisticated queries have also been proposed. A *metric paired comparison* asks the user to (approximately) quantify the difference in utility between two items. A *sorting query* asks the user to sort all items in a collection; it is equivalent to asking all pairwise comparisons between items in the collection. A *bound query* (also referred to as *gamble query*) asks the user to consider a single outcome, and decide whether its value is greater or lower than some specified bound (this can be viewed as a choice between the considered outcome and a random outcome with binomial distribution supported on the best and worst outcomes, with probabilities b and $1-b$, respectively). We refer the reader to e.g., Braziunas and Boutilier (2010) for more detailed information on queries in preference elicitation.

Since only a moderate number of queries can be made, the preferences of the decision-maker are typically not fully understood after the elicitation stage. This is exacerbated by the fact that, when given a choice between alternatives, individuals may sometimes respond in seemingly "*irrational*" ways, being influenced by

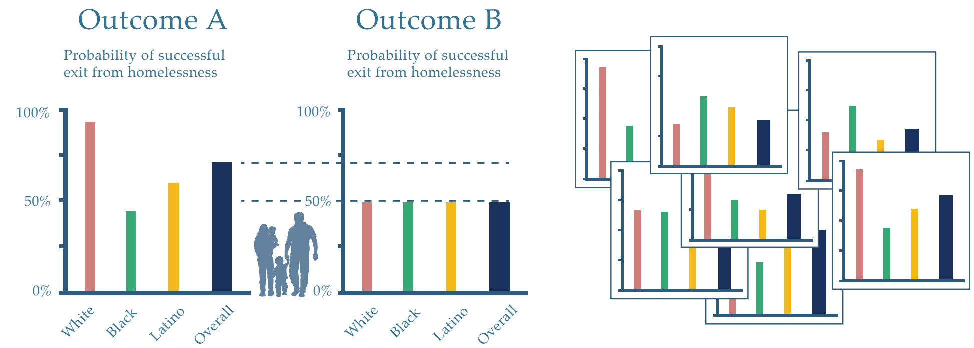

Figure 1 Preference elicitation for housing allocation at LAHSA. The figure on the left shows a very simple sample pairwise comparison: the user is asked whether they would prefer a policy with outcomes A (less fair, more efficient) or B (more fair, less efficient). The figure on the right illustrates the fact that even if one only looks at one dimension of fairness and one dimension of efficiency, there may be many outcomes (candidate policies) to choose from to compare.

the framing of the question, see Allais (1953) and Kahneman and Tversky (1979). Thus, preference elicitation techniques are limited in the information they may provide and any subsequent recommendation needs to be made under *uncertainty* in the utility function of the user. If information about the utility function is encoded by means of a distribution, then the resulting recommendation problem can be formulated as a stochastic program. On the other hand, if utility function information is encoded through membership in a set, the recommendation problem can be expressed as a robust optimization problem which either maximizes worst-case utility (see e.g., Ben-Tal et al. (2009)) or minimizes worst-case regret (see e.g., Savage (1951), Milnor (1954), Stoye (2011)).

1.2. Motivating Applications

In this work, we are motivated by decision-making problems that are hard for human decision-makers to solve, and where preferences over alternatives need to be well understood for an automated decision-making system to be able to make high quality recommendations. Some concrete examples follow.

Public Housing Allocation. The main motivation behind this paper, which we also investigate in our computational experiments, is the problem of allocating scarce housing resources to those experiencing homelessness. Specifically, we aim to build a recommendation system for matching housing resources in Los Angeles County to individuals experiencing homelessness in a way that optimizes the preferences of 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. Currently, there are an estimated 58,000 individuals experiencing homelessness in the L.A. area and only about 22,000 housing resources to help support these individuals. There are two types of housing resources: Permanent Supportive Housing (PSH) and Rapid Rehousing (RRH). Different policies (and associated assignments of resources to individuals) will result in different outcomes in terms of fairness, efficiency, and transparency of the system (e.g., varying wait times and rates of exit from homelessness across the entire population and by race, gender, and sexual orientation). In this case, the recommendation problem is a hard combinatorial problem (a stochastic/robust assignment problem). At the same time, it is hard for homeless service providers to articulate their preferences over such outcomes (attributes) upfront given that hard compromises have to be made. Thus, a preference elicitation step is needed to understand the preferences of policy-makers before making recommendations for matching individuals to resources. In fact, our partners at LAHSA have shared with us that the problem of quantifying preferences in terms of the trade-offs between fairness, efficiency, and interpretability is a recurring theme at their policy council meetings where policy-level decisions are made. In this work, we propose to elicit the policy-maker preferences over policy characteristics (e.g., fairness/efficiency/interpretability trade-off) by asking pairwise comparisons over policy outcomes, see Figure 1 for a simplified example. In our numerical experiments, we use real data from the Homeless Management Information System[4] (HMIS) to show how our proposed algorithm can strategically select a subset of questions to ask the policy-makers to be able to recommend their preferred policy.

Personalized Route Recommendations. In recent years, route recommendation systems have become pervasive. For example, in 2017, Google Maps reported reaching over 1 billion monthly users, see Popper (2017). Different users will typically have different preferences over route attributes (e.g., travel time, travel distance, congestion level, majority of highways or surface streets). Thus, a preference elicitation step is needed to first understand the preferences of drivers before making route recommendations. After preferences have been elicited, the recommendation problem takes the form of a stochastic/robust shortest path or travelling salesman problem which is a hard combinatorial problem. Comparison queries can again be constructed in this setting by putting the user in different situations (e.g., a simplified query may take the following form: "You are late for a meeting. Would you prefer to employ route A with expected travel time 20 minutes and with probability 10% that the travel time will be longer than 40 minutes; or, would you prefer route B with expected travel time 25 minutes and worst-case travel time 30 minutes?").

Budget Allocation. Cities across the U.S., such as the City of L.A. for example, are given a finite budget that they must choose how to allocate among various items[5] (e.g., trees, street sweeping, sidewalks, pothole repair, policing, fire hydrants, encampments, lane painting, illegal dumping, public trash cans, parking permits, and enforcement). Depending on how budget is allocated, outcomes in the city will change (for example, number of robberies, number of complaints about potholes). A recommendation system should aid the city in allocating its funds between various items. In this case, the recommendation problem is a hard stochastic/robust knapsack problem. Comparison queries can be constructed by putting the user in different scenarios about the state of their city, e.g., by showing them statistics on the outcomes.

Note that in the examples above, the choice of a recommendation does not necessarily reduce to choosing the best option among a given set of candidates: entirely *new/custom* products can be designed/synthesized.

1.3. Literature Review

The preference elicitation problem can be viewed as a sequential decision-making problem affected by uncertainty over a finite planning horizon, where *information available* at any point in time is *decision-dependent* (i.e., *endogenous*). At the beginning of each period, a query (or collection of queries) is made and the answer to the query (or queries) made is observed before the next (set of) queries is selected. The choice of queries to make at each period is allowed to *adapt* to the history of queries and their answers. At the end of the planning horizon, an optimal recommendation is made based on the knowledge acquired. Our proposed approach relates to both the literature on preference elicitation and to the works on stochastic and robust optimization with *decision-dependent information discovery.* We review both of these in what follows.

1.3.1. Preference Elicitation Literature Preference elicitation approaches can be classified into one of two frameworks depending on how they model and update uncertainty in the utility function. The first framework takes a stochastic/Bayesian approach with aim to optimize expected utility. The second framework takes a robust/set based viewpoint and usually aims to optimize either worst-case regret or worst-case utility.

Stochastic Methods. In this framework, a *probabilistic prior* is placed over the possible utility functions. This prior typically takes the form of a density function over the utility function parameters. Each time the answer to a query is observed, the prior is updated into a posterior (which acts as the prior for the next period). After all queries are made and all answers are observed, the option with the greatest expected utility is recommended. This model is computationally intractable and thus, to the best of our knowledge, all proposed approaches rely on (greedy) heuristics and approximation schemes to identify queries, see e.g., Chajewska et al. (2000), Boutilier (2002), Brochu et al. (2007), Zhao et al. (2018), and references therein.

Robust/Polyhedral Methods. Our approach most closely relates to the second framework which takes a robust viewpoint by modeling uncertainty in the utility function as deterministic and set based. At the beginning of the planning horizon, an *uncertainty set* of all feasible utilities is built. This set typically takes the form of a polyhedron containing all feasible realizations of the utility function parameters. As queries are made and answered, the uncertainty set is augmented with additional constraints that prune-out any realizations of the utility function that are incompatible with the answers given. At the end of the planning horizon, the option with the highest worst-case utility or lowest worst-case regret is recommended. To the best of our knowledge, the first papers to take this approach originate in the marketing literature. Motivated by the popularity of online conjoint analysis techniques (see e.g., Johnson (1987, 1991)) and by the need for methods that provide reasonable estimates with fewer questions in problems involving many parameters, Toubia et al. (2003, 2004) proposed polyhedral methods for metric paired comparisons and pairwise comparisons, respectively. With polyhedral estimation, at each iteration, the question that is likely to reduce the size of the uncertainty set the fastest is chosen. At the end of the planning horizon, and under the assumption that all elements of the uncertainty set are equally likely, the recommendation is made which performs best when the utility vector is given as the analytic center of the uncertainty polyhedron (to maximize expected utility). Both the elicitation and recommendation tasks rely on efficient mathematical programming techniques and are thus suitable for interactive use. Almost concurrently, Wang and Boutilier (2003) investigate min-max regret based approaches for eliciting user preferences using gamble queries. They show that myopically optimal queries that optimize various improvement criteria can be computed in polynomial time. Boutilier et al. (2006) devise several procedures, based on mixed integer linear optimization, to compute min-max regret solutions to recommendation problems. They also propose a number of heuristic methods for utility elicitation applicable to pairwise comparison and gamble queries. Toubia et al. (2007) provide a probabilistic interpretation of polyhedral methods and propose improvements that incorporate response error and/or informative priors into individual-level question selection and estimation. Bertsimas and O'Hair (2013) investigate pairwise comparison queries and generalize the approach from Toubia et al. (2004) to allow for inconsistencies in user responses. The authors propose to use the max-min utility decision criterion and robust optimization techniques to identify recommendations that are robust to uncertainty in the utility. Several of the above studies demonstrate promising performance of polyhedral methods on both simulated and real preference data (field tests), see e.g., Toubia et al. (2003, 2004) and Braziunas and Boutilier (2010). The approaches of Toubia et al. (2004), Boutilier et al. (2006), and Bertsimas and O'Hair (2013) apply in our context and we will thus benchmark against them in our experiments.

In the present paper, we follow the robust/polyhedral approach to uncertainty modeling. However, we move a significant step beyond the approaches above in that we study *optimal* approaches to preference elicitation that *integrate the elicitation phase with the downstream recommendation* in a single two- or multi-stage robust optimization problem with decision-dependent information discovery.

1.3.2. Literature on Optimization with Decision-Dependent Information Discovery This type of problem was first studied in the context of stochastic programming where it was generally assumed 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 decisions that determine the time of information discovery. This research began with the works of Jonsbråten et al. (1998) and Jonsbråten (1998). More recently, Goel and Grossman (2004) provided mixed-binary programming formulations of scenario based stochastic programs with decision-dependent information discovery. Unfortunately, these formulations are exponential in the number of endogenous uncertain parameters and thus they propose a conservative solution approach that precommits the measurement decisions in the first period. 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. Colvin and Maravelias (2010) and Gupta and Grossmann (2011) investigate iterative solution schemes based on relaxations of the non-anticipativity constraints. Our paper most closely relates to the works of Vayanos et al. (2011) and Vayanos et al. (2019), wherein the authors investigate two- and multi-stage stochastic and robust programs with decision-dependent information discovery that involve continuously supported uncertain parameters. Vayanos et al. (2011) propose a decision-rule based approximation approach that relies on a pre-partitioning of the support of the uncertain parameters. Vayanos et al. (2019) propose a solution method based on the K-adaptability approximation. They also investigate an active preference elicitation problem; their approach however does not apply to the case of comparison queries. In fact, none of the above approaches apply in our setting which presents a combination of discrete and continuous uncertain parameters, namely the responses to the queries and the vector of utility function coefficients.

1.4. Proposed Approach & Contributions

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

(a) We investigate the active preference elicitation problem under the robust/polyhedral approach to uncertainty modelling, under both the max-min utility and min-max regret decision criteria. We propose the first (to the best of our knowledge) formal mathematical formulation of the robust active preference elicitation problem that integrates the learning and recommendation phases under this model. We investigate two settings: *a)* an *offline elicitation* setting, where all queries are made at once, and *b)* an *online elicitation* setting, where queries are selected sequentially over time in an adaptive fashion. The offline (resp. online) active elicitation problem takes the form of a two-stage (resp. multi-stage) robust optimization problem with decision-dependent information discovery involving both discrete and real-valued uncertain parameters.

(b) We show that the offline and online robust active elicitation problems (under both the max-min utility and min-max regret decision criteria) are $\mathcal{NP}$-hard even if the recommendation set consists of only two items. We also demonstrate that, under the max-min utility criterion, if the recommendation set is convex, then there is no value in being strategic about the queries to ask.

(c) In the case of the offline elicitation problem, we provide an enumeration-based algorithm which applies when the number of queries (and items to compare) is small. For larger number of queries, we provide an equivalent reformulation in the form of a mixed-binary linear program. We augment this formulation with symmetry breaking constraints and a column-and-constraint generation algorithm to speed-up computation. For the online elicitation setting, we propose a conservative solution approach combined with a folding horizon strategy. We show that under this approximation, the online active elicitation problem reduces to solving a sequence of offline active elicitation problems.

(d) We perform case studies based on both synthetic data and real data from the HMIS. For the real data case we design policies for prioritizing homeless youth for housing resources in a way that meets the preferences of policy-makers. We demonstrate competitive performance relative to the state of the art in terms of solution time and solution quality. Our case study also highlights the benefits of minimizing worst-case regret relative to maximizing worst-case utility.

1.5. Organization of the Paper and Notation

The remainder of the paper is organized as follows. Section 2 formalizes the model of our recommender system and the preference model that underlies our approach. Sections 3 and 4 study the offline and online problems, respectively. The min-max regret versions of these problems are investigated in Section 5. Section 6 generalizes the approaches in the previous sections to handle inconsistent user responses. Several strategies for speeding-up computation are proposed in Section 7. Section 8 describes our numerical results and Section 9 concludes. 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 $\boldsymbol{x}_k$. We let $\mathbf{e}$ (resp. $\mathbf{e}_i$) denote the vector of all ones (resp. the ith basis vector) of appropriate dimension. 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. 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. Model

In this section, we define *items* (goods) in terms of their attributes, formalize the *set based preference model* that underlies our approach, introduce the notion of a *(comparison) query* that can be used to elicit user preferences, and describe the information gained as a byproduct of an answer to a query. We introduce the robust recommendation problems which, given a set based model of utility uncertainty, recommend a product that either *maximizes* worst-case utility or *minimizes* worst-case regret. We also derive the robust counterparts of these problems. In this way, this section lays the foundations for the computation of *optimal* active preference elicitation strategies in Sections 3, 4, 5, and 6.

2.1. Items (Goods)

In the spirit of multi-attribute utility theory (see Section 1) we assume that, when choosing one item over another, a user is basing their decision on the attributes of the two options. Thus, each item $\boldsymbol{x}$ is uniquely characterized by its J attributes and can therefore be modelled as a point in a J-dimensional Cartesian space, i.e., $\boldsymbol{x} = (\boldsymbol{x}_1, \ldots, \boldsymbol{x}_J)$, where $\boldsymbol{x}_j \in \mathbb{R}$, $j = 1, \ldots, J$, denotes the jth attribute of item $\boldsymbol{x} \in \mathbb{R}^J$. In this work, we are motivated by settings where J may be high-dimensional (taking values in the order of 10 or 20 for example). We denote the universe of all feasible (realizable) items by $\mathcal{X} \subseteq \mathbb{R}^J$ –this set can be thought of as collecting all product configurations that are possible to produce.

Example 1 (Attributes of Policies for Homeless Services Provision). *Policies for matching scarce housing resources to individuals experiencing homelessness can be characterized in terms of various attributes quantifying their fairness, efficiency, and interpretability characteristics, see Vayanos et al. (2019). These include (but are not limited to): expected wait time (overall, by gender, by race), probability of receiving a resource of each type (by gender, by race), probability of exiting homelessness (overall, by gender, by race), number of features used in the policy, depth of the tree for the case of decision-tree-based policies, see e.g., Azizi et al. (2018). In this case, the feasible set $\mathcal{X}$ corresponds to the set of attribute values that can be attained by any given policy in the space of allowable policies (e.g., linear or decision-tree based policies) taking into account resource limitations.*

2.2. User Preferences

We assume that the user has an (unknown) not necessarily strict preference over the items in the universe $\mathcal{X}$. For two items $\boldsymbol{x}$, $\boldsymbol{y} \in \mathcal{X}$, we mark the weak preference relation by $\preceq$ so that $\boldsymbol{x} \preceq \boldsymbol{y}$ means that "the agent prefers $\boldsymbol{y}$ at least as much as $\boldsymbol{x}$" or "the agent weakly prefers $\boldsymbol{y}$ to $\boldsymbol{x}$." Accordingly, we use the symbol $\sim$ as a shorthand to the indifference relation so that $\boldsymbol{x} \sim \boldsymbol{y}$ if and only if $(\boldsymbol{x} \preceq \boldsymbol{y}) \wedge (\boldsymbol{y} \preceq \boldsymbol{x})$, which reads "the agent is indifferent between $\boldsymbol{x}$ and $\boldsymbol{y}$". Lastly, we employ the symbol $\prec$ to indicate the strong preference relation so that $\boldsymbol{x} \prec \boldsymbol{y}$ if and only if $(\boldsymbol{x} \preceq \boldsymbol{y}) \wedge (\boldsymbol{y} \not\preceq \boldsymbol{x})$, which reads "the agent strictly prefers $\boldsymbol{y}$ to $\boldsymbol{x}$."

To streamline presentation, we begin by making the following assumption, common in rational choice theory, see Arrow (1963), and relax it later in Section 6.

Assumption 1 (Rationality). *We assume that the user is rational in the sense that their preferences over items satisfy the following axioms:*

1. *Completeness: For all $\boldsymbol{x}$ and $\boldsymbol{y} \in \mathcal{X}$, we have $\boldsymbol{x} \succ \boldsymbol{y}$, or $\boldsymbol{x} \prec \boldsymbol{y}$, or $\boldsymbol{x} \sim \boldsymbol{y}$.*
2. *Transitivity: For all $\boldsymbol{x}$, $\boldsymbol{y}$, and $\boldsymbol{z} \in \mathcal{X}$, it holds that if item $\boldsymbol{x}$ is weakly preferred to item $\boldsymbol{y}$, and item $\boldsymbol{y}$ is weakly preferred to $\boldsymbol{z}$, then $\boldsymbol{x}$ is weakly preferred to $\boldsymbol{z}$. Mathematically, we have:*

$$\boldsymbol{x} \succeq \boldsymbol{y} \quad \wedge \quad \boldsymbol{y} \succeq \boldsymbol{z} \quad \Rightarrow \quad \boldsymbol{x} \succeq \boldsymbol{z} \quad \forall \boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z} \in \mathcal{X}.$$

3. *Antisymmetry: If $\boldsymbol{x} \succeq \boldsymbol{y}$ and $\boldsymbol{y} \succeq \boldsymbol{x}$, then $\boldsymbol{x} \sim \boldsymbol{y}$.*

Assumption 1 implies that, given the universe $\mathcal{X}$ of exhaustive and exclusive items to choose from, the user can rank the elements of this set in terms of their preferences in a consistent way (the ranking constitutes a total ordering), and the set has at least one maximal element.

We propose to represent user preferences with an (unknown) utility function $u : \mathcal{X} \to \mathbb{R}$ and analyze the user's behavior indirectly with utility functions. Note that, by Debreu's Representation Theorem, this is always possible under Assumption 1 if $\mathcal{X}$ has finite cardinality or if the preference relation is continuous, see Debreu (1954). Thus, u ranks each item in the universe $\mathcal{X}$. The user strictly prefers $\boldsymbol{x}$ to $\boldsymbol{y}$ if and only if $u(\boldsymbol{x}) > u(\boldsymbol{y})$. Accordingly, they are indifferent between $\boldsymbol{x}$ and $\boldsymbol{y}$ if and only if $u(\boldsymbol{x}) = u(\boldsymbol{y})$. When we relax Assumption 1 in Section 6, we will also propose a way to model user preferences via utility functions while directly capturing inconsistencies.

We make the following assumption regarding the utility function.

Assumption 2 (Linear Utility and Polyhedral Uncertainty). *The user utility function $u : \mathcal{X} \to \mathbb{R}$ is linear being expressible as $u(\boldsymbol{x}) := \mathscr{U}^\top \boldsymbol{x}$ for some (random) vector $\mathscr{U}$ supported in the uncertainty set $\mathcal{U}^0 \subseteq \mathbb{R}^J$. Moreover, the set $\mathcal{U}^0$ is a non-empty full-dimensional bounded polyhedron, given by $\mathcal{U}^0 := \{\boldsymbol{u} \in \mathbb{R}^J \mid \boldsymbol{B}\boldsymbol{u} \geq \boldsymbol{b}\}$ for some matrix $\boldsymbol{B} \in \mathbb{R}^{M \times J}$ and vector $\boldsymbol{b} \in \mathbb{R}^M$.*

The assumption above is very common in the literature on preference elicitation, see e.g., Toubia et al. (2003, 2004, 2007), Boutilier et al. (2006), Bertsimas and O'Hair (2013). It is in fact the central assumption behind polyhedral methods in marketing, see Section 1.3.1. The assumption that prior information on $\mathscr{U}$ can be encoded using linear inequality constraints as in the set $\mathcal{U}^0$ is very natural. Indeed, as will become clear later on (see Section 2.4), the type of queries we propose gives rise to such inequalities. Assumption 2 can be relaxed to allow that $\mathcal{U}^0$ presents equality constraints (rather than only inequalities) and that it possesses an inner point (rather than being full-dimensional).

Example 2. *If no prior information is available on the random utility coefficients $\mathscr{U}$, then one may take, without loss of generality, $\mathcal{U}^0 = [-1, 1]^J$, as in e.g., Bertsimas and O'Hair (2013). Indeed, the utility coefficients can be scaled by a constant without affecting the preference ordering over items.*

Example 3. *One may restrict the utility function coefficients to add-up to one as in e.g., Toubia et al. (2003). In that case, $\mathcal{U}^0 = \{\mathbf{u} \in \mathbb{R}_+^J : \mathbf{e}^\top \mathbf{u} = 1\}$ and the coefficients $\boldsymbol{u}$ can be viewed as partworth utilities, i.e., numerical scores that measure how much each feature influences the user's decision to make that choice.*

Example 4. *If item* $\overline{\boldsymbol{x}}$ *is a benchmark, one may normalize utilities relative to that item, by letting* $\mathcal{U}^0 = \{\boldsymbol{u} \in [-1,1]^J \, : \, \boldsymbol{u}^\top \overline{\boldsymbol{x}} = 1\}$.

2.3. Query Set and Recommendation Set

We let $\mathcal{Q} \subseteq \mathcal{X}$ denote the *query set*, i.e., the set of items that the recommender system can use to build queries. We make the following assumption on the query set.

Assumption 3 (Finite Query Set). *The query set* $\mathcal{Q}$ *has finite cardinality* $I \geq 2$. *Items in the query set are indexed by* $i \in \mathcal{I} := \{1, \ldots, I\}$ *and the ith item in the set is denoted by* $\boldsymbol{x}^i \in \mathbb{R}^J$. *Thus,* $\mathcal{Q} = \{\boldsymbol{x}^i \mid i \in \mathcal{I}\}$.

Assumption 3 is very common in the literature. In fact, we have not been able to find a paper that does not make this assumption. At the same time, this assumption is a very natural one. First, in most applications, to be able to compare items, these items must exist. Second, this assumption holds naturally if the support of all attributes is finite (e.g., few configurations for each attribute).

We denote by $\mathcal{R} \subseteq \mathcal{X}$ the *recommendation set*, i.e., the collection of items from which the system can draw to make a recommendation. In general, the sets $\mathcal{Q}$ and $\mathcal{R}$ need not coincide. The ability to cater for cases where $\mathcal{Q}$ and $\mathcal{R}$ are distinct has received little attention in the literature but is important in practical applications. Indeed, we may be able to *design a new item (or product)*, that is not currently available, to meet the needs of a decision- or policy-maker, in which case $\mathcal{R}$ is a strict superset of $\mathcal{Q}$ and may be uncountable, countably infinite, or have a hard combinatorial structure. On the other hand, $\mathcal{R}$ may be a strict subset of $\mathcal{Q}$ if for example some items, that are in principle offered, are out of stock. Next, we describe some examples of recommendation sets that are useful in applications.

Example 5 (Shortest Path Recommendation). *Probably one of the most common decision-support systems consists in recommending routes from a given source* s *to a given destination* t *on a directed graph* $(\mathcal{V}, \mathcal{A})$ *with node set* $\mathcal{V}$ *and edge set* $\mathcal{A}$ *so as to minimize a measure of user disutility, see Section 1.2. In this case, the recommendation set takes the form of the feasible set of a shortest path problem given by*

$$\mathcal{R} = \left\{ \boldsymbol{x} \in \{0,1\}^{|\mathcal{A}|} \, : \, \sum_{j:(i,j)\in\mathcal{A}} \boldsymbol{x}_{ij} - \sum_{j:(j,i)\in\mathcal{A}} \boldsymbol{x}_{ji} = \mathbb{I}(i=s) - \mathbb{I}(i=t) \quad \forall i \in \mathcal{V} \right\}.$$

Example 6 (Knapsack Composition Recommendation). *In numerous applications ranging from budget allocation (see Section 1.2), project portfolio selection, and portfolio optimization, the recommendation set involves a knapsack constraint. This knapsack recommendation set is given by*

$$\mathcal{R} = \left\{ \boldsymbol{x} \in \{0,1\}^J \, : \, \sum_{j=1}^{J} \boldsymbol{w}_j \boldsymbol{x}_j \, \leq W \right\},$$

where $\boldsymbol{w}_j \in \mathbb{R}_+$, $j = 1, \ldots, J$, represent weights associated with undertaking project/investment j (associated with choice $\boldsymbol{x}_j = 1$) and $W \in \mathbb{R}_+$ represents a budget limit. This set consists of a single constraint that requires that the collection of projects/investments undertaken does not consume more than the available budget W.

Example 7 (Assignment Recommendation). *A variety of recommendation problems are restricted by assignment constraints. These arise when a number of agents indexed in the set $\mathcal{A}$ need to be assigned to a number of tasks indexed in the set $\mathcal{T}$, incurring some cost that may vary depending on the agent-task assignment. It is required to perform all tasks by assigning each of them to an agent. The assignment recommendation set takes the form*

$$\mathcal{R} = \left\{ \boldsymbol{x} \in \{0,1\}^{|\mathcal{A}| \times |\mathcal{T}|} \, : \, \sum_{i \in \mathcal{A}} \boldsymbol{x}_{ij} = 1 \;\; \forall j \in \mathcal{T}, \quad \sum_{j \in \mathcal{T}} \boldsymbol{x}_{ij} = 1 \;\; \forall i \in \mathcal{A} \right\}.$$

It arises in numerous applications such as the matching of customers to cars in ridesharing apps (e.g., Uber[6] or Lyft[7]), or in the allocation of scarce resources such as kidneys (see e.g., Bertsimas et al. (2013), Bandi et al. (2018)) or public housing (see e.g., Azizi et al. (2018), Vayanos et al. (2019)).

Example 8 (Minimum Spanning Tree Recommendation). *A variety of practical recommendation problems that arise in applications are restricted by spanning tree constraints. These arise in the design of networks, including computer networks, telecommunication networks, transportation networks, water supply networks, and electrical grids. Given a directed graph $(\mathcal{V}, \mathcal{A})$ with node set $\mathcal{V}$ and edge set $\mathcal{A}$, this recommendation set takes the form*

$$\mathcal{R} = \left\{ \boldsymbol{x} \in \{0,1\}^{|\mathcal{A}|} \, : \, \sum_{(i,j) \in \mathcal{A}} \boldsymbol{x}_{ij} \, = |\mathcal{V}| - 1, \quad \sum_{(i,j) \in \mathcal{A} : i \in \mathcal{S}, j \in \mathcal{S}} \boldsymbol{x}_{ij} \leq |\mathcal{S}| - 1 \quad \forall \mathcal{S} \subseteq \mathcal{V} \right\}.$$

2.4. Elicitation through Comparison Queries

Before recommending an item (or outcome) from the set $\mathcal{R}$, the recommender system has the opportunity to make a number of queries to the user. These queries (may) enable the system to gain information about $\mathscr{U}$ (see Assumption 2) thus improving the quality of the recommendation. Each query takes the form of a comparison between two items in the query set. We denote the set of all comparisons the system may choose from by $\mathcal{C} := \{(i,i') \mid i \in \mathcal{I},\, i' \in \mathcal{I},\, i < i'\}$. Note that only comparisons of items i and i' for $i < i'$ need to be considered since the order of the items in each comparison does not matter. The set $\mathcal{C}$ can be augmented with constraints that preclude for example items that are very similar from being compared. Such constraints may be useful in practical settings where preferences are elicited using graphs, where minute differences in feature values may not be distinguishable to the naked eye. Similarly, the set $\mathcal{C}$ can include constraints that ensure that items to be compared only differ in a small number of features. Again, such constraints are important in practical settings, where it may be hard, from a cognitive standpoint, for users to compare high-dimensional items. We omit such constraints to minimize notational overheard.

We allow the recommender system to make K queries, indexed in the set $\mathcal{K} := \{1,\ldots,K\}$, before making a recommendation. We let $\boldsymbol{\iota}^\kappa := (\boldsymbol{\iota}_1^\kappa, \boldsymbol{\iota}_2^\kappa) \in \mathcal{C}$ denote the κth query, $\kappa \in \mathcal{K}$. Thus, $\boldsymbol{\iota}_1^\kappa$ and $\boldsymbol{\iota}_2^\kappa$ denote the indices of the first and second items in the κth query, respectively: Query κ asks the user to compare $\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa}$ and $\boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}$. In particular, the user must choose one of three possible answers in response to query κ: *(a)* "I strictly prefer $\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa}$ to $\boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}$" (i.e., $\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} \succ \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}$); *(b)* "I am indifferent between $\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa}$ and $\boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}$" (i.e., $\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} \sim \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}$); or *(c)* "I strictly prefer $\boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}$ to $\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa}$" (i.e., $\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} \prec \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}$).

We associate each possible answer to query κ to a response scenario $\boldsymbol{s}_\kappa \in \mathcal{S} := \{-1, 0, 1\}$ such that

$$\boldsymbol{s}_\kappa = \begin{cases} 1 & \text{if } \boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} \succ \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa} \\ 0 & \text{if } \boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} \sim \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa} \\ -1 & \text{else.} \end{cases}$$

The information obtained on $\mathscr{U}$ depends on the answer to query κ, i.e., on the response scenario. Each choice imposes a different linear constraint on the random utility coefficients $\mathscr{U}$, see Assumption 2: *(a)* If $\boldsymbol{s}_\kappa = 1$, then $\mathscr{U}^\top \boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} > \mathscr{U}^\top \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}$; *(b)* If $\boldsymbol{s}_\kappa = 0$, then $\mathscr{U}^\top \boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} = \mathscr{U}^\top \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}$; and *(c)* If $\boldsymbol{s}_\kappa = -1$, then $\mathscr{U}^\top \boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} < \mathscr{U}^\top \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}$. We collect the answers to each query in the vector $\boldsymbol{s} := \{\boldsymbol{s}_\kappa\}_{\kappa \in \mathcal{K}}$ and accordingly let $\boldsymbol{\iota} := \{(\boldsymbol{\iota}_1^\kappa, \boldsymbol{\iota}_2^\kappa)\}_{\kappa \in \mathcal{K}}$. After all K queries have been made and the responses to the queries are observed, we can update the support of $\mathscr{U}$ (i.e., the set of all possible realizations of $\mathscr{U}$) as follows (see Assumption 2):

$$\mathcal{U}(\boldsymbol{\iota}, \boldsymbol{s}) = \left\{ \begin{array}{rl} \boldsymbol{u} \in \mathcal{U}^0 \,:\, \boldsymbol{u}^\top(\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}) > 0 & \forall \kappa \in \mathcal{K} \,:\, \boldsymbol{s}_\kappa = 1 \\ \boldsymbol{u}^\top(\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}) = 0 & \forall \kappa \in \mathcal{K} \,:\, \boldsymbol{s}_\kappa = 0 \\ \boldsymbol{u}^\top(\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}) < 0 & \forall \kappa \in \mathcal{K} \,:\, \boldsymbol{s}_\kappa = -1 \end{array} \right\}. \tag{1}$$

We emphasize that the support of $\mathscr{U}$ after the K comparisons depends both on $\boldsymbol{\iota}$, the queries made, and on $\boldsymbol{s}$, the answers given by the user to the queries, see Figure 2.

2.5. Risk Averse & Regret Averse Recommendations

After K queries have been made and the answers to these queries have been observed, the recommendation system needs to select (or design) an item from the (possibly uncountable, countably infinite, or combinatorial) set $\mathcal{R}$ to recommend to the user. At the time when the recommendation is made, the coefficients $\mathscr{U}$ are still unknown: they are merely known to belong to the uncertainty set $\mathcal{U}(\boldsymbol{\iota}, \boldsymbol{s})$. It is thus natural that the system seeks to provide recommendations that are *robust* to all possible realization of $\mathscr{U}$ in the set $\mathcal{U}(\boldsymbol{\iota}, \boldsymbol{s})$. We investigate two notions of robustness that are popular in the literature: recommendations that either maximize worst-case utility (see e.g., Bertsimas and O'Hair (2013)) or that minimize worst-case regret (see e.g., Boutilier et al. (2006)).

Maximizing Worst-Case Utility. Given uncertainty in the utility function coefficients, it is natural for risk averse decision-makers to seek recommendations that will maximize the worst-case utility of the recommended item. Mathematically, given the sequences, $\boldsymbol{\iota}$ and $\boldsymbol{s}$ of questions and answers, the recommender system offers the item with the maximum worst-case (minimum) utility for any $\boldsymbol{u} \in \mathcal{U}(\boldsymbol{\iota}, \boldsymbol{s})$. Such a *risk averse recommendation* solves the problem

$$\underset{\boldsymbol{x} \in \mathcal{R}}{\text{maximize}} \quad \min_{\boldsymbol{u} \in \mathcal{U}(\boldsymbol{\iota}, \boldsymbol{s})} \quad \boldsymbol{u}^\top \boldsymbol{x}. \tag{$\mathcal{R}_{\text{risk}}$}$$

Note in particular that the optimal solution to this problem depends on both $\boldsymbol{\iota}$ and $\boldsymbol{s}$. While the recommender system has no control over the agent responses, $\boldsymbol{s}$, it *can* select the comparisons $\boldsymbol{\iota} \in \mathcal{C}^K$ so as to improve the quality of the recommendation.

Minimizing Worst-Case Regret. Given uncertainty in the utility function coefficients, certain decision-makers exhibit regret (rather than risk) aversion: they anticipate regret and thus incorporate in their choice their desire to reduce it. 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

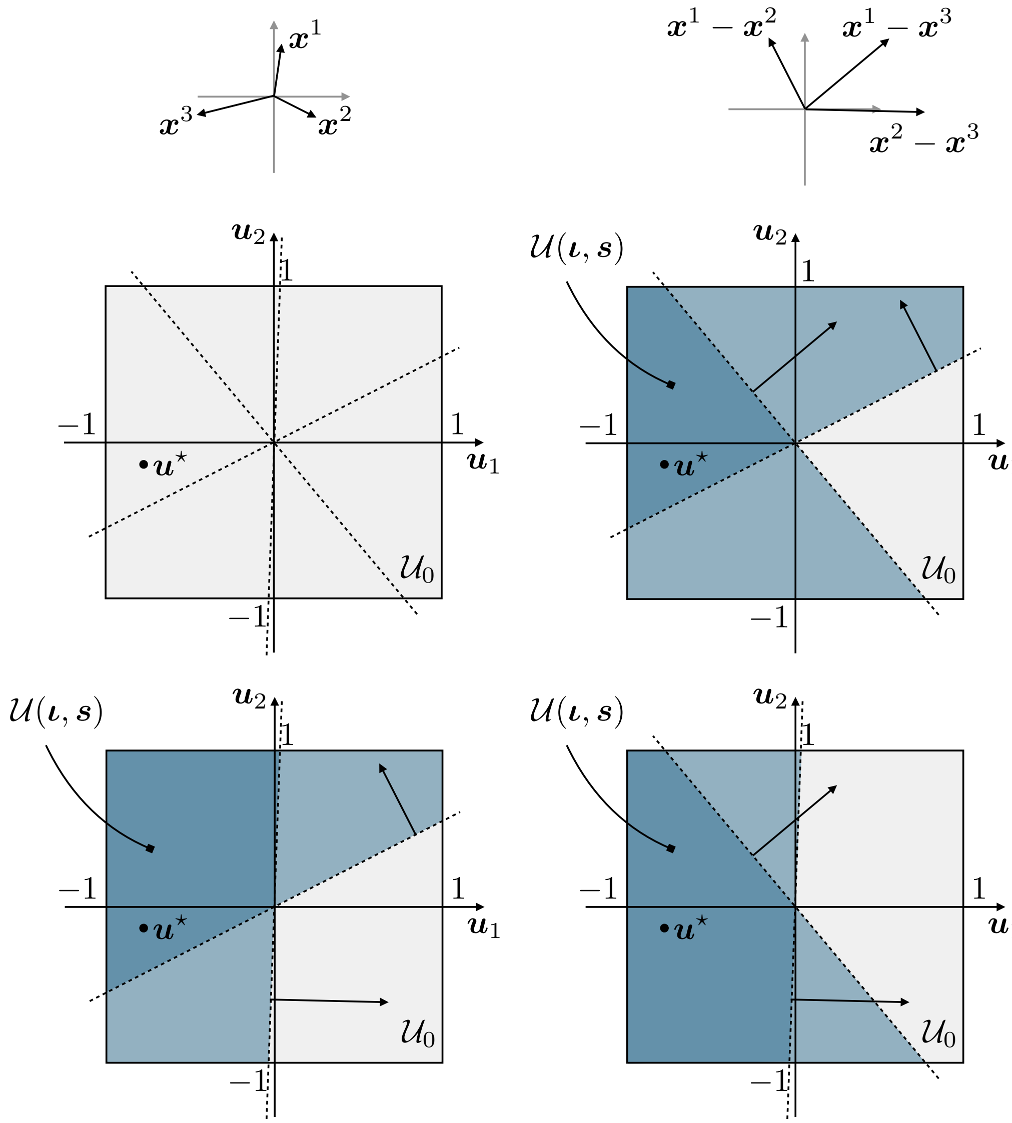

Figure 2 Illustration of the uncertainty set update procedure when there are $I = 3$ items, each with $J = 2$ features, the "true" (but unknown) utility vector is $\boldsymbol{u}^\star$, and the system is allowed to ask $K = 2$ questions before making a recommendation. The first row shows the three items $\boldsymbol{x}^1$, $\boldsymbol{x}^2$, and $\boldsymbol{x}^3$ in $\mathbb{R}^2$ (L) and the vectors $\boldsymbol{x}^1 - \boldsymbol{x}^2$, $\boldsymbol{x}^1 - \boldsymbol{x}^3$, and $\boldsymbol{x}^2 - \boldsymbol{x}^3$ associated with each of the comparison queries $(\boldsymbol{x}^1, \boldsymbol{x}^2)$, $(\boldsymbol{x}^1, \boldsymbol{x}^3)$, and $(\boldsymbol{x}^2, \boldsymbol{x}^3)$, respectively (R). The left figure on the second row shows the initial uncertainty set $\mathcal{U}^0$, the vector $\boldsymbol{u}^\star$, and the hyperplanes associated with each of the queries. The remaining three figures show the uncertainty set $\mathcal{U}(\boldsymbol{\iota}, \boldsymbol{s})$ updated in response to the three different pairs of queries $\{(\boldsymbol{x}^1, \boldsymbol{x}^2), (\boldsymbol{x}^1, \boldsymbol{x}^3)\}$ (row 2, R), $\{(\boldsymbol{x}^1, \boldsymbol{x}^2), (\boldsymbol{x}^2, \boldsymbol{x}^3)\}$ (row 3, L), and $\{(\boldsymbol{x}^1, \boldsymbol{x}^3), (\boldsymbol{x}^2, \boldsymbol{x}^3)\}$ (row 3, R). Note that the uncertainty set changes depending on the queries asked and on the answers given, which in turn depend on the underlying vector $\boldsymbol{u}^\star$ (whose value is unknown by the recommender system).

in practical applications. Given the sequences, $\boldsymbol{\iota}$ and $\boldsymbol{s}$, the recommender system offers the item with the minimum worst-case (maximum) regret for any $\boldsymbol{u} \in \mathcal{U}(\boldsymbol{\iota}, \boldsymbol{s})$. Such a *regret averse recommendation* solves the problem

$$\underset{\boldsymbol{x}\in\mathcal{R}}{\text{minimize}} \quad \max_{\boldsymbol{u}\in\mathcal{U}(\boldsymbol{\iota},\boldsymbol{s})} \quad \left\{ \max_{\boldsymbol{x}'\in\mathcal{R}} \; \boldsymbol{u}^\top \boldsymbol{x}' - \boldsymbol{u}^\top \boldsymbol{x} \right\}. \tag{$\mathcal{R}_{\text{regret}}$}$$

As in the risk averse case, the optimal solution to this problem depends on both $\boldsymbol{\iota}$ and $\boldsymbol{s}$. The choice of comparison queries $\boldsymbol{\iota} \in \mathcal{C}^K$ should thus be guided by the downstream recommendation problem.

To the best of our knowledge, all preference elicitation techniques from the literature are heuristic and do not provide guarantees on performance. In the present paper, we propose to compute *provably optimal* queries that optimize worst-case utility (see Sections 3 and 4) and worst-case regret (see Section 5). In the remainder of this paper, we focus on selecting a set of comparisons $\boldsymbol{\iota}$ that maximizes (resp. minimizes) the objective of Problem ($\mathcal{R}_{\text{risk}}$) (resp. ($\mathcal{R}_{\text{regret}}$)) when the responses to the queries are adversarially chosen in the sense that they are as "*uninformative*" as possible: they hinder the optimal values of the downstream recommendation problems as much as possible.

2.6. Two Active Preference Elicitation Strategies: Offline and Online Elicitation

In this paper, we investigate two strategies for eliciting the preferences of the user. In the first, termed *offline preference elicitation*, K queries are selected in advance, before any answer is revealed. We investigate offline preference elicitation with the max-min utility (resp. min-max regret) decision criterion in Section 3 (resp. Sections 5.1-5.4). In the second, termed *online preference elicitation*, K queries are selected one at a time and the answer to each query is revealed before the next query is selected. We study online preference elicitation with the max-min utility (resp. min-max regret) decision criterion in Section 4 (resp. Section 5.5).

Offline elicitation is for example common in low-resource settings where interactions with the user do not involve a computer (e.g., paper-based questionnaires) or if it is necessary that all users be presented with the same queries (e.g., in a controlled study). Online elicitation on the other hand is preferred in settings where a computer is available and where different users can be presented with different options. Indeed, by allowing the questions to *adjust* (or *adapt*) to the user responses, better questions can be chosen, yielding a "smaller" uncertainty set $\mathcal{U}(\boldsymbol{\iota}, \boldsymbol{s})$ and a less conservative recommendation.

We emphasize that both the offline and online elicitation strategies are *active* in the sense that the choice of queries to make is directly informed by the downstream recommendation problem: different recommendation sets will yield different choices of queries. We believe that this sets us apart from existing literature.

3. Offline Active Elicitation with the Max-Min Utility Decision Criterion

We begin by addressing the *offline active preference elicitation problem with the max-min utility decision criterion*, which we also refer to as the *risk averse offline preference elicitation* problem. In this problem, all K comparisons, $\boldsymbol{\iota} := \{\boldsymbol{\iota}^\kappa\}_{\kappa\in\mathcal{K}}$, are precommitted in advance, before any agent responses are observed. After agent responses are observed, the uncertainty set is updated and an item with highest worst-case utility is recommended. We have two motivations for studying this problem. First, the offline problem where all K comparisons are selected in advance is interesting in its own right, see Section 2.6. Second, the offline problem is a useful building block for (approximately) solving the online preference elicitation problem, see Section 4. This section is organized as follows. In Section 3.1, we formulate the offline preference elicitation problem as a max-min-max-min optimization problem and study its complexity. We develop an enumeration based solution approach and an equivalent reformulation in the form of a mixed-binary linear program (MBLP) that can be solved with off-the-shelf solvers in Sections 3.2 and 3.3, respectively. Finally, in Section 3.4, we propose a column-and-constraint generation approach for solving the MBLP.

3.1. Problem Formulation & Complexity Analysis

In the risk averse offline preference elicitation problem, the sequence of events is as follows. First, the recommender system selects K queries $\boldsymbol{\iota}^\kappa \in \mathcal{C}$, $\kappa \in \mathcal{K}$, to ask the user. Subsequently, the user, who has a true (but unknown) utility vector $\boldsymbol{u}^\star$ from $\mathcal{U}^0$, responds to the queries truthfully and rationally (see Assumption 1) by selecting answers $\boldsymbol{s}_\kappa$ to each query $\boldsymbol{\iota}^\kappa$ in a way that complies with their utility vector $\boldsymbol{u}^\star$. Note that the utility vector $\boldsymbol{u}^\star$ is not observable to the recommender system; only the answers to the questions are. In fact, the user themself is not aware of their vector $\boldsymbol{u}^\star$ (else, they would directly share it with the recommender system). Once the answers to the queries are observed, the recommender system can certify that $\boldsymbol{u}^\star \in \mathcal{U}(\boldsymbol{\iota}, \boldsymbol{s})$ and solves the risk averse recommendation problem ($\mathcal{R}_{\text{risk}}$) by offering an item that is robust to all utility vectors in $\mathcal{U}(\boldsymbol{\iota}, \boldsymbol{s})$. In this section, since the decision-maker is risk averse, it is natural that they be hedged against adversarial responses $\boldsymbol{s}$.

Mathematically, the risk averse offline preference elicitation problem is expressible as the following two-stage robust optimization problem with decision-dependent information discovery

$$\underset{\boldsymbol{\iota}\in\mathcal{C}^K}{\text{maximize}} \quad \min_{\boldsymbol{s}\in\mathcal{S}(\boldsymbol{\iota})} \quad \max_{\boldsymbol{x}\in\mathcal{R}} \quad \min_{\boldsymbol{u}\in\mathcal{U}(\boldsymbol{\iota},\boldsymbol{s})} \quad \boldsymbol{u}^\top \boldsymbol{x}, \tag{$\mathcal{P}^K_{\text{off,risk}}$}$$

where $\mathcal{S}(\boldsymbol{\iota}) := \{\boldsymbol{s} \in \mathcal{S}^K \mid \mathcal{U}(\boldsymbol{\iota}, \boldsymbol{s}) \neq \emptyset\}$ denotes the set of all answers compatible with some $\boldsymbol{u} \in \mathcal{U}^0$. Indeed, the set $\mathcal{U}(\boldsymbol{\iota}, \boldsymbol{s})$ is empty if and only if either the utility vector used to generate the answers is not in the set $\mathcal{U}^0$ or if at least one of the answers provided was not rational, in the sense that it did not comply with the chosen $\boldsymbol{u} \in \mathcal{U}^0$. Since the set $\mathcal{U}^0$ is non-empty by Assumption 2, the set $\mathcal{S}(\boldsymbol{\iota})$ is by construction non-empty for all $\boldsymbol{\iota} \in \mathcal{C}^K$. Problem ($\mathcal{P}^K_{\text{off,risk}}$) is a "two-and-a-half" stage robust optimization problem with decision-dependent uncertainty set.

Remark 1. Single-stage robust *optimization problems with decision-dependent uncertainty sets have been investigated by Spacey et al. (2012), Nohadani and Sharma (2016), Nohadani and Roy (2017), Zhang et al. (2017), and Lappas and Gounaris (2018). Similarly,* single-stage distributionally robust *optimization problems have been studied by Noyan et al. (2018), Basciftci et al. (2019), Luo and Mehrotra (2019), and Ryu and Jiang (2019). To the best of our knowledge, the only papers to investigate* two- *and* multi-stage *problems with decision-dependent uncertainty are Bertsimas and Vayanos (2017) and Vayanos et al. (2011, 2019). The papers Vayanos et al. (2011, 2019) are the only ones to apply to the case of decision-dependent information discovery. Neither of these approaches applies in our context which presents a combination of discretely and continuously supported uncertain parameters. At the same time, approaches for solving two-stage robust optimization problems (see e.g., Bertsimas and Dunn (2017), Bertsimas and Georghiou (2015, 2018), Bertsimas and Dunning (2016), Postek and Den Hertog (2016), Bertsimas and Caramanis (2010), Hanasusanto et al. (2015), Subramanyam et al. (2017), Chassein et al. (2019), Rahmattalabi et al. (2019)) do not apply to Problem ($\mathcal{P}^K_{\text{on,risk}}$), which presents decision-dependent information discovery and discretely supported uncertain parameters. Thus, new solution techniques are required.*

Problem ($\mathcal{P}^K_{\text{off,risk}}$) is difficult to solve for many reasons. First, it is a max-min-max-min problem. Second, the uncertainty sets $\mathcal{S}(\boldsymbol{\iota})$ and $\mathcal{U}(\boldsymbol{\iota}, \boldsymbol{s})$ for the first and second decision-stages are decision-dependent. Third, the uncertainty set $\mathcal{U}(\boldsymbol{\iota}, \boldsymbol{s})$ depends upon the first stage uncertainty $\boldsymbol{s}$ and is also open, making it difficult to derive computational solution approaches. The following lemma shows that Problem ($\mathcal{P}^K_{\text{off,risk}}$) can be considerably simplified by eliminating the dependence of $\mathcal{S}(\boldsymbol{\iota})$ on $\boldsymbol{\iota}$ and by replacing the strict inequalities in the set $\mathcal{U}(\boldsymbol{\iota}, \boldsymbol{s})$ by their loose counterparts.

Lemma 1. *Problem ($\mathcal{P}^K_{\text{off,risk}}$) is equivalent to*

$$\underset{\boldsymbol{\iota} \in \mathcal{C}^K}{\text{maximize}} \quad \min_{\boldsymbol{s} \in \mathcal{S}^K} \quad \max_{\boldsymbol{x} \in \mathcal{R}} \quad \min_{\boldsymbol{u} \in \widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s})} \quad \boldsymbol{u}^\top \boldsymbol{x}, \tag{2}$$

where

$$\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s}) := \left\{ \begin{array}{rl} \boldsymbol{u}\in\mathcal{U}^0 \,:\, \boldsymbol{u}^\top(\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa}-\boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}) \geq 0 & \forall\kappa\in\mathcal{K}\,:\,\boldsymbol{s}_\kappa=1 \\ \boldsymbol{u}^\top(\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa}-\boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}) = 0 & \forall\kappa\in\mathcal{K}\,:\,\boldsymbol{s}_\kappa=0 \\ \boldsymbol{u}^\top(\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa}-\boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}) \leq 0 & \forall\kappa\in\mathcal{K}\,:\,\boldsymbol{s}_\kappa=-1 \end{array} \right\}$$

in the sense that the two problems have the same optimal objective value and the same sets of optimal solutions.

Lemma 1 follows naturally. First, note that since $\mathcal{S}(\boldsymbol{\iota})$ is non-empty, there exist $\boldsymbol{s}\in\mathcal{S}(\boldsymbol{\iota})$ such that $\mathcal{U}(\boldsymbol{\iota},\boldsymbol{s})\neq\emptyset$, implying that the optimal objective value of Problem ($\mathcal{P}^K_{\text{off,risk}}$) is finite. On the other hand, any choice of $\boldsymbol{s}\in\mathcal{S}^K\backslash\mathcal{S}(\boldsymbol{\iota})$ will result in the inner maximization over $\boldsymbol{u}$ to be taken over an empty set, resulting in an objective equal to $+\infty$, implying that this choice of $\boldsymbol{s}$ is suboptimal. We can thus include such choices of $\boldsymbol{s}$ in the first minimization without modifying the optimal value of the problem. Second, using Assumption 2 and classical results in polyhedral theory, it can be shown that for any given $\boldsymbol{s}\in\mathcal{S}^K$ such that $\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s})$ is non-empty, we can construct $\boldsymbol{s}'\in\mathcal{S}^K$ such that $\mathcal{U}(\boldsymbol{\iota},\boldsymbol{s}')\neq\emptyset$ and $\text{cl}(\mathcal{U}(\boldsymbol{\iota},\boldsymbol{s}'))=\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s})$. This implies that we can relax the strict inequalities in the uncertainty set and retain an equivalent problem.

Before studying the complexity of the risk averse offline preference elicitation problem, we show that the number of response scenarios in Problem (2) can be drastically reduced.

Observation 1. *Problem* (2) *is equivalent to*

$$\underset{\boldsymbol{\iota}\in\mathcal{C}^K}{\text{maximize}} \quad \min_{\boldsymbol{s}\in\widetilde{\mathcal{S}}^K} \quad \max_{\boldsymbol{x}\in\mathcal{R}} \quad \min_{\boldsymbol{u}\in\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s})} \quad \boldsymbol{u}^\top\boldsymbol{x}, \tag{$\widetilde{\mathcal{P}}^K_{\text{off,risk}}$}$$

where $\widetilde{\mathcal{S}}:=\{-1,1\}$, *in the sense that the two problems have the same optimal objective value and the same sets of optimal solutions.*

Observation 1 is very natural. It states that it is never in "nature's" favor to select scenarios $\boldsymbol{s}$ such that $\boldsymbol{s}_\kappa=0$ for some $\kappa\in\{1,\ldots,K\}$. Indeed, declaring indifference as the answer to any query in Problem (2) always results in an objective value that is no lower than that obtained by considering "$\succeq$" and "$\preceq$" responses only.

Next, we show that Problem ($\widetilde{\mathcal{P}}^K_{\text{off,risk}}$) is $\mathcal{NP}$-hard when the recommendation set is discrete, which motivates the integer programming based reformulations we provide in the next sections.

Theorem 1. *The following claims about the complexity of Problem* ($\widetilde{\mathcal{P}}^K_{\text{off,risk}}$) *hold true.*

(a) Suppose the recommendation set $\mathcal{R}$ *is convex. Then, Problem* ($\widetilde{\mathcal{P}}^K_{\text{off,risk}}$) *is polynomially solvable. Moreover, it holds that Problems* ($\widetilde{\mathcal{P}}^K_{\text{off,risk}}$) *and* ($\widetilde{\mathcal{P}}^0_{\text{off,risk}}$) *are equivalent for all* $K \in \mathbb{N}$, *i.e., in the case of polyhedral recommendation set, in the worst-case, there is no benefit in asking any queries. In addition, the optimal objective values of Problems* ($\widetilde{\mathcal{P}}^K_{\text{off,risk}}$) *and* ($\mathcal{R}_{\text{risk}}$) *coincide.*

(b) Evaluating the objective function of Problem ($\widetilde{\mathcal{P}}^K_{\text{off,risk}}$) *is* $\mathcal{NP}$*-hard even if the recommendation set* $\mathcal{R}$ *consists of only two elements.*

Motivated by the complexity results above, we henceforth assume the recommendation set in Problem ($\widetilde{\mathcal{P}}^K_{\text{off,risk}}$) is discrete and MBLP representable. We provide an enumeration-based approach for solving ($\widetilde{\mathcal{P}}^K_{\text{off,risk}}$) in Section 3.2, which is motivated by settings where K is moderately valued. We provide a general, MBLP reformulation of Problem ($\widetilde{\mathcal{P}}^K_{\text{off,risk}}$) in Section 3.3 and a column-and-constraint generation approach in Section 3.4. These are motivated by settings where K is large.

3.2. An Enumeration-Based Solution Approach for Small K

In the case of a small number of queries K and moderate number of items I, our approach relies on the following proposition which shows that for any fixed $\boldsymbol{\iota} \in \mathcal{C}$ and $\boldsymbol{s} \in \widetilde{\mathcal{S}}^K$, the robust recommendation problem ($\mathcal{R}_{\text{risk}}$) is equivalent to a single maximization problem of size polynomial in the size of the input.

Proposition 1. *For any fixed* $\boldsymbol{\iota} \in \mathcal{C}^K$ *and* $\boldsymbol{s} \in \widetilde{\mathcal{S}}^K$, *the robust recommendation problem* ($\mathcal{R}_{\text{risk}}$) *is equivalent to the maximization problem*

$$\begin{array}{ll} \text{maximize} & \boldsymbol{b}^\top \boldsymbol{\beta} \\ \text{subject to} & \boldsymbol{x} \in \mathcal{R},\ \boldsymbol{\alpha} \in \mathbb{R}^K_+,\ \boldsymbol{\beta} \in \mathbb{R}^M_+ \\ & \displaystyle\sum_{\kappa \in \mathcal{K}} \boldsymbol{s}_\kappa \left(\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa} \right) \boldsymbol{\alpha}_\kappa + \boldsymbol{B}^\top \boldsymbol{\beta} = \boldsymbol{x}, \end{array} \tag{3}$$

whose size is polynomial in the size of the input. Problem (3) *is a mixed-binary linear program.*

The proof of this result relies on classical robust optimization techniques, see e.g., Ben-Tal et al. (2009) and Gorissen et al. (2015). Our enumeration-based algorithm proposes to evaluate the objective value of Problem (3) for all choices of $\boldsymbol{\iota} \in \mathcal{C}^K$ and $\boldsymbol{s} \in \widetilde{\mathcal{S}}^K$ to determine the choice of $\boldsymbol{\iota}$ that will yield the highest worst-case objective. The total number of queries whose performance we need to evaluate is $|\mathcal{C}|^K = (|\mathcal{Q}|(|\mathcal{Q}|-1))^K$ and the total number of scenarios is $|\widetilde{\mathcal{S}}^K| = 2^K$. Thus, for small values of K (i.e., in the order of 1 to 3) and

---

Algorithm 1: Enumeration-based algorithm for solving $(\widetilde{\mathcal{P}}^K_{\text{off,risk}})$ for moderate values of K.

Input: Comparison set $\mathcal{C}$, recommendation set $\mathcal{R}$, uncertainty set $\mathcal{U}^0$, number of queries K;

Output: Optimal query $\boldsymbol{\iota}^\star$ from $\mathcal{C}^K$;

Initialization: $\boldsymbol{\iota}^\star \leftarrow \emptyset$; OPT $\leftarrow -\infty$;

for $\boldsymbol{\iota} \in \mathcal{C}^K$ **do**

  foreach $\boldsymbol{s} \in \widetilde{\mathcal{S}}^K$ **do**

    $\text{OPT}(\boldsymbol{\iota}, \boldsymbol{s}) \leftarrow$ optimal objective of Problem (3);

  end

  if $\min_{\boldsymbol{s} \in \widetilde{\mathcal{S}}^K} \text{OPT}(\boldsymbol{\iota}, \boldsymbol{s}) > \text{OPT}$ **then**

    $\text{OPT} \leftarrow \min_{\boldsymbol{s} \in \widetilde{\mathcal{S}}^K} \text{OPT}(\boldsymbol{\iota}, \boldsymbol{s})$; $\boldsymbol{\iota}^\star \leftarrow \boldsymbol{\iota}$;

  end

end

---

moderate values of I (i.e., in the order of 10 to 20), it is computationally practicable to enumerate all choices of $\boldsymbol{\iota} \in \mathcal{C}^K$ and $\boldsymbol{s} \in \widetilde{\mathcal{S}}^K$.

We now propose to leverage Proposition 1 and the moderate cardinality of $\mathcal{C}^K$ and $\widetilde{\mathcal{S}}^K$ to devise an enumeration-based algorithm for solving Problem $(\widetilde{\mathcal{P}}^K_{\text{off,risk}})$. The idea is as follows: for each $\boldsymbol{\iota} \in \mathcal{C}^K$, we solve 2^K instances of Problem (3) (one for each element $\boldsymbol{s}$ of $\widetilde{\mathcal{S}}^K$) and record the minimum value of the objective over these instances. The optimal solution to Problem $(\widetilde{\mathcal{P}}^K_{\text{off,risk}})$ is given by the choice of $\boldsymbol{\iota}$ that gives the greatest minimum value. The total number of problems solved is given by $(2|\mathcal{C}|)^K$ and can thus be done efficiently if K is small. In particular, if $K=1$, it reduces to solving $2|\mathcal{C}|$ MBLPs. For clarity, our approach is detailed in Algorithm 1. We note that although Algorithm 1 is very simple and intuitive in nature, it has not (to the best of our knowledge) been proposed in the literature.

3.3. Exact MBLP Reformulation

In most real world applications, K is in the order of 8 to 10. In such settings, the number of possible query combinations (i.e., the cardinality of $\mathcal{C}^K$) and response scenarios (i.e., the cardinality of $\widetilde{\mathcal{S}}^K$) are both large so that the enumeration approach proposed in Section 3.2 becomes computationally prohibitive. In this section, we propose an exact MBLP reformulation to Problem ($\widetilde{\mathcal{P}}^K_{\text{off,risk}}$), which enables us to select queries optimally even when K is large. This reformulation can leverage techniques from integer optimization to speed-up computation and circumvent complete enumeration.

The first step towards reformulating Problem ($\widetilde{\mathcal{P}}^K_{\text{off,risk}}$) consists in noting that we can exchange the order of the inner minimization and maximization problems, provided we allow the choice of recommended item to depend on the response scenario $\boldsymbol{s}$. This statement is formalized in the following observation.

Observation 2. *Problem ($\widetilde{\mathcal{P}}^K_{\text{off,risk}}$) is equivalent to the following max-min problem*

$$\underset{\boldsymbol{\iota}\in\mathcal{C}^K}{\text{maximize}} \quad \max_{\substack{\boldsymbol{x}^{\boldsymbol{s}}\in\mathcal{R}:\\ \boldsymbol{s}\in\widetilde{\mathcal{S}}^K}} \quad \min_{\boldsymbol{s}\in\widetilde{\mathcal{S}}^K} \quad \min_{\boldsymbol{u}\in\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s})} \quad \boldsymbol{u}^\top \boldsymbol{x}^{\boldsymbol{s}}, \tag{4}$$

where $\boldsymbol{x}^{\boldsymbol{s}}$ denotes the item to recommend in response scenario $\boldsymbol{s}$, $\boldsymbol{s}\in\widetilde{\mathcal{S}}^K$.

Observation 2 is the key to convert the *"two-and-a-half" stage* (min-max-min-max) Problem (4) to a *single-stage* robust problem as shown in the following lemma.

Lemma 2. *Problem* (4) *is equivalent to the following finite program*

$$\begin{array}{ll}
\text{maximize} & \tau \\
\text{subject to} & \tau\in\mathbb{R},\ \boldsymbol{\iota}\in\mathcal{C}^K \\
& \left.\begin{array}{l}
\boldsymbol{\alpha}^{\boldsymbol{s}}\in\mathbb{R}^K_+,\ \boldsymbol{\beta}^{\boldsymbol{s}}\in\mathbb{R}^M_+,\ \boldsymbol{x}^{\boldsymbol{s}}\in\mathcal{R} \\
\tau \ \le\ \boldsymbol{b}^\top\boldsymbol{\beta}^{\boldsymbol{s}} \\
\displaystyle\sum_{\kappa\in\mathcal{K}} \boldsymbol{s}_\kappa\,(\boldsymbol{x}^{\boldsymbol{\iota}^\kappa_1}-\boldsymbol{x}^{\boldsymbol{\iota}^\kappa_2})\,\boldsymbol{\alpha}^{\boldsymbol{s}}_\kappa + \boldsymbol{B}^\top\boldsymbol{\beta}^{\boldsymbol{s}} \ =\ \boldsymbol{x}^{\boldsymbol{s}}
\end{array}\right\} \quad \forall \boldsymbol{s}\in\widetilde{\mathcal{S}}^K,
\end{array} \tag{5}$$

where $\boldsymbol{\iota}$ denotes the queries to make and $\boldsymbol{x}^{\boldsymbol{s}}$ the items to recommend in response scenario $\boldsymbol{s}\in\widetilde{\mathcal{S}}^K$.

Next, we convert Problem (5) to a finite MBLP by introducing, for each query, binary variables which indicate which item is included as first and second element in the query. Specifically, we encode the choice of a query $\boldsymbol{\iota}^\kappa\in\mathcal{C}$, $\kappa\in\mathcal{K}$, using two sets of binary decision variables, $\boldsymbol{v}^\kappa\in\{0,1\}^I$ and $\boldsymbol{w}^\kappa\in\{0,1\}^I$, whose ith element is one if and only if item i is the first (resp. second) item in query κ. Equivalently, $\boldsymbol{\iota}^\kappa=(i,i')$ if and only if $\boldsymbol{v}^\kappa_i=\boldsymbol{w}^\kappa_{i'}=1$. The following theorem shows that Problem (5) can be reformulated as an MBLP.

Theorem 2. *Problem* (5) *is equivalent to the following mixed-binary linear program*

$$
\begin{array}{lll}
\text{maximize} & \tau & \\
\text{subject to} & \tau \in \mathbb{R},\, \boldsymbol{v}^\kappa,\, \boldsymbol{w}^\kappa \in \{0,1\}^I,\, \kappa \in \mathcal{K} & \\
& \boldsymbol{\alpha}^{\boldsymbol{s}} \in \mathbb{R}^K_+,\, \boldsymbol{\beta}^{\boldsymbol{s}} \in \mathbb{R}^M_+,\, \boldsymbol{x}^{\boldsymbol{s}} \in \mathcal{R},\, \boldsymbol{s} \in \widetilde{\mathcal{S}}^K & \\
& \overline{\boldsymbol{v}}^{\boldsymbol{s}\kappa},\, \overline{\boldsymbol{w}}^{\boldsymbol{s}\kappa} \in \mathbb{R}^I_+,\, \boldsymbol{s} \in \widetilde{\mathcal{S}}^K,\, \kappa \in \mathcal{K} & \\
& \left.\begin{array}{l} \tau \;\leq\; \boldsymbol{b}^\top \boldsymbol{\beta}^{\boldsymbol{s}} \\ \displaystyle\sum_{i \in \mathcal{I}} \boldsymbol{x}^i \sum_{\kappa \in \mathcal{K}} \boldsymbol{s}_\kappa \left(\overline{\boldsymbol{v}}^{\boldsymbol{s}\kappa}_i - \overline{\boldsymbol{w}}^{\boldsymbol{s}\kappa}_i \right) + \boldsymbol{B}^\top \boldsymbol{\beta}^{\boldsymbol{s}} \;=\; \boldsymbol{x}^{\boldsymbol{s}} \end{array}\right\} & \forall \boldsymbol{s} \in \widetilde{\mathcal{S}}^K \\
& \mathbf{e}^\top \boldsymbol{v}^\kappa = 1,\, \mathbf{e}^\top \boldsymbol{w}^\kappa = 1,\, 1 - \boldsymbol{w}^\kappa_i \;\geq\; \displaystyle\sum_{i' : i' \geq i} \boldsymbol{v}^\kappa_{i'} \qquad \forall i \in \mathcal{I},\, \forall \kappa \in \mathcal{K} & \\
& \left.\begin{array}{l} \overline{\boldsymbol{v}}^{\boldsymbol{s}\kappa} \;\leq\; M \boldsymbol{v}^\kappa,\; \overline{\boldsymbol{v}}^{\boldsymbol{s}\kappa} \;\leq\; \boldsymbol{\alpha}^{\boldsymbol{s}}_\kappa \mathbf{e},\; \overline{\boldsymbol{v}}^{\boldsymbol{s}\kappa} \;\geq\; \boldsymbol{\alpha}^{\boldsymbol{s}}_\kappa \mathbf{e} - M(\mathbf{e} - \boldsymbol{v}^\kappa) \\ \overline{\boldsymbol{w}}^{\boldsymbol{s}\kappa} \;\leq\; M \boldsymbol{w}^\kappa,\; \overline{\boldsymbol{w}}^{\boldsymbol{s}\kappa} \;\leq\; \boldsymbol{\alpha}^{\boldsymbol{s}}_\kappa \mathbf{e},\; \overline{\boldsymbol{w}}^{\boldsymbol{s}\kappa} \;\geq\; \boldsymbol{\alpha}^{\boldsymbol{s}}_\kappa \mathbf{e} - M(\mathbf{e} - \boldsymbol{w}^\kappa) \end{array}\right\} & \begin{array}{c} \forall \boldsymbol{s} \in \widetilde{\mathcal{S}}^K \\ \kappa \in \mathcal{K}, \end{array}
\end{array}
\tag{6}
$$

where M is a "big-M" constant. In particular, given an optimal solution $(\tau, \{\boldsymbol{v}^\kappa, \boldsymbol{w}^\kappa\}_{\kappa \in \mathcal{K}}, \{\boldsymbol{\alpha}^{\boldsymbol{s}}, \boldsymbol{\beta}^{\boldsymbol{s}}, \boldsymbol{x}^{\boldsymbol{s}}\}_{\boldsymbol{s} \in \widetilde{\mathcal{S}}^K})$ to Problem (6)*, an optimal set of queries for the risk averse offline preference elicitation problem is given by*

$$
\boldsymbol{\iota}^\kappa_1 = \sum_{i \in \mathcal{I}} i \cdot \mathbb{I}\left(\boldsymbol{v}^\kappa_i = 1\right) \qquad \textit{and} \qquad \boldsymbol{\iota}^\kappa_2 = \sum_{i \in \mathcal{I}} i \cdot \mathbb{I}\left(\boldsymbol{w}^\kappa_i = 1\right), \quad \kappa \in \mathcal{K}.
$$

The first and second sets of constraints in Problem (6) originate directly from formulation (5). The third set of constraints guarantees that the κth query $\boldsymbol{\iota}^\kappa$ is such that $\boldsymbol{\iota}^\kappa_1 < \boldsymbol{\iota}^\kappa_2$, i.e., $(\boldsymbol{\iota}^\kappa_1, \boldsymbol{\iota}^\kappa_2) \in \mathcal{C}$. The fourth and fifth sets of constraints enable us to linearize the products $\boldsymbol{\alpha}^{\boldsymbol{s}}_\kappa \boldsymbol{v}^\kappa$ and $\boldsymbol{\alpha}^{\boldsymbol{s}}_\kappa \boldsymbol{w}^\kappa$.

Reformulation (6) is very attractive as it enables us to solve the risk averse offline preference elicitation problem $(\mathcal{P}^K_{\text{off,risk}})$ as an MBLP using off-the-shelf solvers. For fixed K, Problem (6) is polynomial in the size of the input. Yet, Problem (6) is exponential in K. Thus, further strategies are needed to ensure that we can solve it faster as K grows. We note that offline surveys will typically involve a moderate numbers of queries (e.g., $K \approx 10$) to avoid tiring the user. Moreover, offline surveys are typically prepared in advance, before interacting with the user, and so do not require instantaneous solution. Compared to the approach proposed in Section 3.2, the MBLP formulation (6) is expected to be more tractable, since it can exploit integer optimization technology to avoid enumerating all possible choices. To ensure that it can scale to practical values of K, in the remainder of this section, we propose decomposition strategies to speed-up solution. Additional strategies based on e.g., strengthening of the formulation are discussed in Section 7.

3.4. Column-and-Constraint Generation

Problem (6) presents a number of decision variables and constraints that are exponential in K, making it impracticable to solve when the number of queries is large. At the same time, due to the robust nature of the problem, we expect that only a moderate number of scenarios $\boldsymbol{s} \in \widetilde{\mathcal{S}}^K$ will be candidates to be active in the epigraph constraint. For this reason, in this section, we propose a column-and-constraint generation algorithm inspired from Vayanos et al. (2019) to speed-up computation and which applies when $\mathcal{R}$ has fixed finite cardinality (in which case the problem is $\mathcal{NP}$-hard, see Theorem 1). To minimize notational overhead, we describe our column-and-constraint generation procedure using the finite program (5). Naturally, all problems solved would need to be converted to MBLPs first using techniques similar to those employed in Theorem 2. We omit these conversions to streamline presentation.

The key idea behind our algorithm is to decompose the problem into a relaxed master problem and a series of subproblems indexed by $\boldsymbol{s} \in \widetilde{\mathcal{S}}^K$. The master problem initially only involves a subset of the constraints (those indexed by $\boldsymbol{s} \in \mathcal{S}' \subseteq \widetilde{\mathcal{S}}^K$) and a *single auxiliary MBLP* is used to iteratively identify indices $\boldsymbol{s} \in \widetilde{\mathcal{S}}^K$ for which the solution to the relaxed master problem becomes infeasible when plugged into subproblem $\boldsymbol{s}$. 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 relaxed master problem parameterized by the index set $\mathcal{S}'$ as

$$\begin{array}{ll} \text{maximize} & \tau \\ \text{subject to} & \tau \in \mathbb{R},\ \boldsymbol{\iota} \in \mathcal{C}^K \\ & \left.\begin{array}{l} \boldsymbol{\alpha}^{\boldsymbol{s}} \in \mathbb{R}^K_+,\ \boldsymbol{\beta}^{\boldsymbol{s}} \in \mathbb{R}^M_+,\ \boldsymbol{x}^{\boldsymbol{s}} \in \mathcal{R} \\ \tau \ \leq \ \boldsymbol{b}^\top \boldsymbol{\beta}^{\boldsymbol{s}} \\ \displaystyle\sum_{\kappa \in \mathcal{K}} \boldsymbol{s}_\kappa \left(\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}\right) \boldsymbol{\alpha}^{\boldsymbol{s}}_\kappa + \boldsymbol{B}^\top \boldsymbol{\beta}^{\boldsymbol{s}} \ = \ \boldsymbol{x}^{\boldsymbol{s}} \end{array}\right\} \quad \forall \boldsymbol{s} \in \mathcal{S}'. \end{array} \qquad (\mathcal{CCG}^{\text{master}}_{\text{risk}}(\mathcal{S}'))$$

Note that this problem only involves a subset of the decision variables and constraints of Problem (5). Given variables $(\tau, \boldsymbol{\iota})$ feasible in the master problem, we define the $\boldsymbol{s}$th subproblem, $\boldsymbol{s} \in \widetilde{\mathcal{S}}^K$, through

$$\begin{array}{ll} \text{maximize} & 0 \\ \text{subject to} & \boldsymbol{\alpha} \in \mathbb{R}^K_+,\ \boldsymbol{\beta} \in \mathbb{R}^M_+,\ \boldsymbol{x} \in \mathcal{R} \\ & \tau \leq \boldsymbol{b}^\top \boldsymbol{\beta} \\ & \displaystyle\sum_{\kappa \in \mathcal{K}} \boldsymbol{s}_\kappa \left(\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}\right) \boldsymbol{\alpha}_\kappa + \boldsymbol{B}^\top \boldsymbol{\beta} \ = \ \boldsymbol{x}. \end{array} \qquad (\mathcal{CCG}^{\text{sub},\boldsymbol{s}}_{\text{risk}}(\tau, \boldsymbol{\iota}))$$

An inspection of the Proof of Proposition 1 reveals that the equality constraint in Problem ($\mathcal{CCG}^{\text{sub},\boldsymbol{s}}_{\text{risk}}(\tau,\boldsymbol{\iota})$) combined with the non-negativity constraints on $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ define the feasible set of the dual of a linear program that is feasible and bounded. The objective function of this dual is $\boldsymbol{b}^\top\boldsymbol{\beta}$. Thus, for τ sufficiently small, Problem ($\mathcal{CCG}^{\text{sub},\boldsymbol{s}}_{\text{risk}}(\tau,\boldsymbol{\iota})$) will be feasible. To identify indices of subproblems ($\mathcal{CCG}^{\text{sub},\boldsymbol{s}}_{\text{risk}}(\tau,\boldsymbol{\iota})$) that, given a solution $(\tau,\boldsymbol{\iota})$ to the relaxed master problem, are infeasible, we solve a *single* feasibility MBLP defined through

$$
\begin{array}{lll}
\text{minimize} & \theta & \\
\text{subject to} & \theta\in\mathbb{R},\ \boldsymbol{s}\in\widetilde{\mathcal{S}}^K,\ \boldsymbol{u}^{\boldsymbol{x}}\in\mathcal{U}^0 \quad \forall \boldsymbol{x}\in\mathcal{R} & \\
& \theta \geq (\boldsymbol{u}^{\boldsymbol{x}})^\top\boldsymbol{x} \quad \forall \boldsymbol{x}\in\mathcal{R} & \qquad (\mathcal{CCG}^{\text{feas}}_{\text{risk}}(\boldsymbol{\iota})) \\
& \left.\begin{array}{l} (\boldsymbol{u}^{\boldsymbol{x}})^\top(\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa}-\boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}) \geq -M(1-\boldsymbol{s}_\kappa) \\ (\boldsymbol{u}^{\boldsymbol{x}})^\top(\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa}-\boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}) \leq M(\boldsymbol{s}_\kappa+1) \end{array}\right\} \quad \forall\kappa\in\mathcal{K},\ \boldsymbol{x}\in\mathcal{R}. &
\end{array}
$$

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

Proposition 2. *Let $\boldsymbol{\iota}$ be feasible in the relaxed master problem ($\mathcal{CCG}^{\text{master}}_{\text{risk}}(\mathcal{S}')$). Then, $\boldsymbol{\iota}$ is feasible in Problem ($\widetilde{\mathcal{P}}^K_{\text{off,risk}}$) and the objective value of $\boldsymbol{\iota}$ in Problem ($\widetilde{\mathcal{P}}^K_{\text{off,risk}}$) is given by the optimal objective value of Problem ($\mathcal{CCG}^{\text{feas}}_{\text{risk}}(\boldsymbol{\iota})$). If $\mathcal{R}$ has fixed finite cardinality, then Problem ($\mathcal{CCG}^{\text{master}}_{\text{risk}}(\mathcal{S}')$) can be solved as an MBLP using off-the-shelf solvers.*

Proposition 2 implies that, for any $\boldsymbol{\iota}$ feasible in the relaxed master problem ($\mathcal{CCG}^{\text{master}}_{\text{risk}}(\mathcal{S}')$), the optimal value of ($\mathcal{CCG}^{\text{feas}}_{\text{risk}}(\boldsymbol{\iota})$) yields a lower bound to the optimal value of Problem ($\widetilde{\mathcal{P}}^K_{\text{off,risk}}$). At the same time, it is evident that for any index set $\mathcal{S}'\subseteq\widetilde{\mathcal{S}}^K$, the optimal value of Problem ($\mathcal{CCG}^{\text{master}}_{\text{risk}}(\mathcal{S}')$) yields an upper bound to the optimal objective value of Problem ($\widetilde{\mathcal{P}}^K_{\text{off,risk}}$). The lemma below is key to identify indices of subproblems $\boldsymbol{s}\in\widetilde{\mathcal{S}}^K$ that are infeasible.

Lemma 3. *The relaxed master problem ($\mathcal{CCG}^{\text{master}}_{\text{risk}}(\mathcal{S}')$) is always feasible. If ($\mathcal{CCG}^{\text{master}}_{\text{risk}}(\mathcal{S}')$) is solvable, let $(\tau,\boldsymbol{\iota},\{\boldsymbol{\alpha}^{\boldsymbol{s}},\boldsymbol{\beta}^{\boldsymbol{s}},\boldsymbol{x}^{\boldsymbol{s}}\}_{\boldsymbol{s}\in\widetilde{\mathcal{S}}})$ be an optimal solution. Else, if ($\mathcal{CCG}^{\text{master}}_{\text{risk}}(\mathcal{S}')$) is unbounded, set $\tau=\infty$ and let $\boldsymbol{\iota}\in\mathcal{C}^K$ be such that ($\mathcal{CCG}^{\text{master}}_{\text{risk}}(\mathcal{S}')$) is unbounded when $\boldsymbol{\iota}$ is fixed to that value. Moreover, let $(\theta,\{\boldsymbol{u}^{\boldsymbol{x}}\}_{\boldsymbol{x}\in\mathcal{R}},\boldsymbol{s})$ be optimal in Problem ($\mathcal{CCG}^{\text{feas}}_{\text{risk}}(\boldsymbol{\iota})$). Then, the following hold:*

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

(ii) If $\theta=\tau$, then Problem ($\mathcal{CCG}^{\text{sub},\boldsymbol{s}}_{\text{risk}}(\tau,\boldsymbol{\iota})$) is feasible for all $\boldsymbol{s}\in\widetilde{\mathcal{S}}^K$;

Algorithm 2: Column-and-Constraint Generation procedure for solving Problem ($\widetilde{\mathcal{P}}^K_{\text{off,risk}}$).

Inputs: Optimality tolerance δ, comparison set $\mathcal{C}$, and recommendation set $\mathcal{R}$;

Initial uncertainty set $\mathcal{U}^0$ and number of queries K;

Output: Query $\boldsymbol{\iota}^\star$ from $\mathcal{C}^K$, near optimal in Problem ($\widetilde{\mathcal{P}}^K_{\text{off,risk}}$) with associated objective θ;

Initialization: $\boldsymbol{\iota}^\star \leftarrow \emptyset$; Upper and lower bounds: $\text{UB} \leftarrow +\infty$ and $\text{LB} \leftarrow -\infty$;

Initialize index set: $\mathcal{S}' \leftarrow \emptyset$;

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

  Solve the master problem ($\mathcal{CCG}^{\text{master}}_{\text{risk}}(\mathcal{S}')$). If it is solvable, let $(\tau, \boldsymbol{\iota}, \{\boldsymbol{\alpha}^{\boldsymbol{s}}, \boldsymbol{\beta}^{\boldsymbol{s}}, \boldsymbol{x}^{\boldsymbol{s}}\}_{\boldsymbol{s}\in\mathcal{S}'})$ be an optimal solution. If it is unbounded, set $\tau = \infty$ and let $\boldsymbol{\iota} \in \mathcal{C}^K$ be such that ($\mathcal{CCG}^{\text{master}}_{\text{risk}}(\mathcal{S}')$) is unbounded when $\boldsymbol{\iota}$ is fixed to that value;

  Set $\text{UB} \leftarrow \tau$;

  Solve the feasibility subproblem ($\mathcal{CCG}^{\text{feas}}_{\text{risk}}(\boldsymbol{\iota})$). Let $(\theta, \{\boldsymbol{u}^{\boldsymbol{x}}\}_{\boldsymbol{x}\in\mathcal{R}}, \boldsymbol{s})$ denote an optimal solution;

  Set $\text{LB} \leftarrow \theta$;

  if $\theta < \tau$ **then**

    $\mathcal{S}' \leftarrow \mathcal{S}' \cup \{\boldsymbol{s}\}$;

  end

end

Set $\boldsymbol{\iota}^\star \leftarrow \boldsymbol{\iota}$;

Result: Collection of queries $\boldsymbol{\iota}^\star$ near-optimal in ($\widetilde{\mathcal{P}}^K_{\text{off,risk}}$) with objective value θ.

(iii) If $\theta < \tau$, then scenario $\boldsymbol{s}$ corresponds to an infeasible subproblem, i.e., Problem ($\mathcal{CCG}^{\text{sub},\boldsymbol{s}}_{\text{risk}}(\tau, \boldsymbol{\iota})$) is infeasible.

Remark 2. *A solution $\boldsymbol{\iota} \in \mathcal{C}^K$ such that ($\mathcal{CCG}^{\text{master}}_{\text{risk}}(\mathcal{S}')$) is unbounded when $\boldsymbol{\iota}$ is fixed to that value can be readily obtained by augmenting formulation ($\mathcal{CCG}^{\text{master}}_{\text{risk}}(\mathcal{S}')$) with an artificial (large) upper bound on τ and solving for an optimal query in that restricted problem.*

Propositions 2 and Lemma 3 culminate in Algorithm 2 whose convergence is guaranteed by the following theorem.

Theorem 3. *Algorithm 2 terminates in a final number of steps with a feasible solution to Problem ($\widetilde{\mathcal{P}}^K_{\text{off,risk}}$). The objective value attained by this solution is within δ of the optimal objective value of the problem.*

In Section 8, we leverage Algorithm 2 and Theorem 3 to solve an active preference learning problem that seeks to recommend housing allocation policies with highest possible worst-case utility.

4. Online Active Elicitation with the Max-Min Utility Decision Criterion

In Section 3, we assumed that all queries are chosen at once. However, in many settings of practical interest, queries are made one at a time and the answer to each query is revealed before the next query is selected. In this case, the recommender system has the opportunity to *adjust* their choice of queries, taking into account the information acquired as a byproduct of previous queries and their answers. In this section, we address this *online active preference elicitation problem with the max-min utility decision criterion* where comparisons are selected adaptively over time. We also refer to this problem as the *risk averse online active preference elicitation problem.*

4.1. Problem Formulation & Complexity Analysis

In the online risk averse recommendation problem, the sequence of events is as follows. First, the recommender system selects K queries $\boldsymbol{\iota}_\kappa \in \mathcal{C}^K$, $\kappa \in \mathcal{K}$, *one at a time*. Each time a query $\boldsymbol{\iota}_\kappa$ is made, and before the next query $\boldsymbol{\iota}_{\kappa+1}$ is selected, the user selects an answer $\boldsymbol{s}_\kappa \in \mathcal{S}$ to query κ which the recommender system observes. As in Section 3, we assume that the user is truthful and rational, see Assumption 1. We relax this assumption later in Section 6. Thus, the sequence of answers given by the user must comply with at least one element $\boldsymbol{u}$ from $\mathcal{U}^0$. We will relax this assumption later in Section 6. After having observed the answer to the first κ queries, the recommender system can assert that the utility vector $\mathscr{U}$ of the user lies in the set

$$\mathcal{U}^\kappa\left(\boldsymbol{\iota}^{[\kappa]}, \boldsymbol{s}_{[\kappa]}\right) := \left\{ \begin{array}{rll} \boldsymbol{u} \in \mathcal{U}^0 \,:\, \boldsymbol{u}^\top(\boldsymbol{x}^{\iota_1^k} - \boldsymbol{x}^{\iota_2^k}) > 0 & \forall k \in \{1, \ldots, \kappa\} \,:\, \boldsymbol{s}_k = 1 \\ \boldsymbol{u}^\top(\boldsymbol{x}^{\iota_1^k} - \boldsymbol{x}^{\iota_2^k}) = 0 & \forall k \in \{1, \ldots, \kappa\} \,:\, \boldsymbol{s}_k = 0 \\ \boldsymbol{u}^\top(\boldsymbol{x}^{\iota_1^k} - \boldsymbol{x}^{\iota_2^k}) < 0 & \forall k \in \{1, \ldots, \kappa\} \,:\, \boldsymbol{s}_k = -1 \end{array} \right\},$$

where $\boldsymbol{\iota}^{[\kappa]} := \{\boldsymbol{\iota}^k\}_{k=1}^\kappa$ and $\boldsymbol{s}_{[\kappa]} := \{\boldsymbol{s}_k\}_{k=1}^\kappa$ denote the history of the first κ queries and answers, respectively. Second, after the sequence of K queries and answers is complete, the recommender systems solves the robust recommendation problem ($\mathcal{R}_{\text{risk}}$) and offers an item from $\mathcal{R}$ that is robust to all utility vectors in $\mathcal{U}^K\left(\boldsymbol{\iota}^{[K]}, \boldsymbol{s}_{[K]}\right) = \mathcal{U}(\boldsymbol{\iota}, \boldsymbol{s})$.

Mathematically, the online risk averse recommendation problem is expressible as

$$\max_{\boldsymbol{\iota}^1 \in \mathcal{C}} \quad \min_{\boldsymbol{s}_1 \in \mathcal{S}^1(\boldsymbol{\iota}^1)} \quad \max_{\boldsymbol{\iota}^2 \in \mathcal{C}} \quad \min_{\boldsymbol{s}_2 \in \mathcal{S}^2(\boldsymbol{\iota}^{[2]}, \boldsymbol{s}_{[1]})} \quad \cdots \quad \max_{\boldsymbol{\iota}^K \in \mathcal{C}} \quad \min_{\boldsymbol{s}_K \in \mathcal{S}^K(\boldsymbol{\iota}^{[K]}, \boldsymbol{s}_{[K-1]})} \quad \max_{\boldsymbol{x} \in \mathcal{R}} \quad \min_{\boldsymbol{u} \in \mathcal{U}(\boldsymbol{\iota}, \boldsymbol{s})} \quad \boldsymbol{u}^\top \boldsymbol{x}, \qquad (\mathcal{P}^K_{\text{on,risk}})$$

where

$$\mathcal{S}^1(\boldsymbol{\iota}^1) := \left\{ \boldsymbol{s}_1 \in \mathcal{S} \mid \mathcal{U}^1(\boldsymbol{\iota}^1, \boldsymbol{s}_1) \neq \emptyset \right\} \quad \text{and} \quad \mathcal{S}^\kappa(\boldsymbol{\iota}^{[\kappa]}, \boldsymbol{s}_{[\kappa-1]}) := \left\{ \boldsymbol{s}_\kappa \in \mathcal{S} \mid \mathcal{U}^\kappa(\boldsymbol{\iota}^{[\kappa]}, \boldsymbol{s}_{[\kappa]}) \neq \emptyset \right\}$$

for each $\kappa \in \mathcal{K}$. Thus, $\mathcal{S}^1(\boldsymbol{\iota}^1)$ denotes the set of all answers to the first query that are compatible with some $\boldsymbol{u} \in \mathcal{U}^0$. Accordingly, $\mathcal{S}^\kappa(\boldsymbol{\iota}^{[\kappa]}, \boldsymbol{s}_{[\kappa-1]})$ denotes the set of all answers to query $\boldsymbol{\iota}_\kappa$ compatible with some $\boldsymbol{u} \in \mathcal{U}^0$ and with the answers given to the first $\kappa - 1$ queries.

The online problem $(\mathcal{P}^K_{\text{on,risk}})$ appears significantly more complicated to solve than its offline counterpart, Problem $(\mathcal{P}^K_{\text{off,risk}})$, due to the $K + 1$ alternating max-min problems. The following theorem shows that, similar to its offline counterpart, the online problem is generally $\mathcal{NP}$-hard.

Theorem 4. *Evaluating the objective function of Problem $(\mathcal{P}^K_{\text{on,risk}})$ is $\mathcal{NP}$-hard even if the sequence of queries is fixed and static and the recommendation set $\mathcal{R}$ consists of only two elements.*

4.2. Conservative Solution Approach: Constant Decision Rule and Folding Horizon

Problem $(\mathcal{P}^K_{\text{on,risk}})$ can be reformulated *equivalently* as an MBLP by performing a sequence of interchanges of max and min operators, resulting in an equivalent max-min-max-min problem, and subsequently using techniques similar to those in Section 3.3. In the max-min-max-min formulation, the queries for any time $\kappa \in \mathcal{K}$ are indexed by $\boldsymbol{s}_{[\kappa-1]}$ and become decision variables of the outermost maximization problem. Unfortunately, the resulting formulation is computationally prohibitive to solve due to the exponential number of decision variables and constraints. Thus, in this section, we propose to follow a *conservative* solution approach based on a *decision rule approximation*, in the spirit of modern robust optimization, see e.g., Ben-Tal et al. (2009), Bertsimas et al. (2010), Gorissen et al. (2015).

To approximately solve the *online* risk averse preference elicitation problem $(\mathcal{P}^K_{\text{on,risk}})$, we propose to solve a sequence of *offline* risk averse preference elicitation problems of the form $(\widetilde{\mathcal{P}}^\kappa_{\text{off,risk}})$, $\kappa = K, \ldots, 1$, in a folding horizon fashion. Thus, for fixed $\kappa \in \mathcal{K}$ and given the sequence of queries and answers $\boldsymbol{\iota}^{[\kappa-1]}$ and $\boldsymbol{s}_{[\kappa-1]}$, the κth query is selected randomly among the set of $K - \kappa + 1$ queries that are optimal to make in the offline preference elicitation problem with uncertainty set $\mathcal{U}^{\kappa-1}(\boldsymbol{\iota}^{[\kappa-1]}, \boldsymbol{s}_{[\kappa-1]})$. Once a query $\boldsymbol{\iota}^\kappa$ is made and the answer $\boldsymbol{s}_\kappa$ to the query is observed, the uncertainty set is updated and a new instance of the offline problem is solved with a "smaller" planning horizon and "smaller" uncertainty set. The process is repeated until K queries have been made at which point an item is offered to the user that is optimal in Problem ($\mathcal{R}_{\text{risk}}$).

Note that to select a query at each period, we are effectively approximating the query selection decisions by *constant* decision rules, while we are still allowing the recommendation made to be *fully adjustable*. We thus obtain a conservative approximation to Problem ($\mathcal{P}^K_{\text{on,risk}}$). The idea of approximating adjustable variables by decision rules of benign complexity is very popular in the stochastic and robust optimization communities, see e.g., Ben-Tal et al. (2004), Kuhn et al. (2009), Bertsimas et al. (2011), Bertsimas and Goyal (2012), Vayanos et al. (2012), Zhen et al. (2016), Xu and Burer (2018), Bodur and Luedtke (2018). As we will see in our numerical experiments, see Section 8, although this approximation is conservative, it significantly outperforms the state of the art approaches from the literature. This is not surprising: our choice of queries is informed by the structure of the downstream recommendation problem. This is in sharp contrast to existing approaches from the literature, where the queries are chosen based on an "information gain" criterion and the choice of query is not informed by the structure of the set $\mathcal{R}$.

5. Active Elicitation with the Min-Max Regret Decision Criterion

In Sections 3 and 4, we assumed that the decision-maker is risk averse. In this section, we instead take the point of view of a *regret averse* decision-maker. We address the *offline active preference elicitation problem with the min-max regret decision criterion*, which we also refer to as the *regret averse offline active preference elicitation problem*, in Sections 5.1-5.4. We study the *online regret averse preference elicitation problem* in Section 5.5.

5.1. Active Offline Elicitation with the Min-Max Regret Decision Criterion

In the regret averse offline active preference elicitation problem, the sequence of events is as follows. First, the recommender system selects K queries $\boldsymbol{\iota}^\kappa \in \mathcal{C}$, $\kappa \in \mathcal{K}$, to ask the user. Subsequently, the user, who has a true (but unknown) utility vector $\boldsymbol{u}^\star$ from $\mathcal{U}^0$, responds to the queries truthfully and rationally (see Assumption 1) by selecting answers $\boldsymbol{s}_\kappa$ to each query $\boldsymbol{\iota}^\kappa$ in a way that complies with their utility vector $\boldsymbol{u}^\star$. As in the risk averse case, the utility vector $\boldsymbol{u}^\star$ is not observable to the recommender system; only the answers to the questions are. Once the answers to the queries are observed, the recommender system can certify that $\boldsymbol{u}^\star \in \mathcal{U}(\boldsymbol{\iota}, \boldsymbol{s})$ and solves the regret averse recommendation problem ($\mathcal{R}_{\text{regret}}$).

Mathematically, the regret averse offline preference elicitation problem is expressible as the following two-stage robust optimization problem with decision-dependent information discovery

$$\underset{\boldsymbol{\iota}\in\mathcal{C}^K}{\text{minimize}} \quad \max_{\boldsymbol{s}\in\mathcal{S}(\boldsymbol{\iota})} \quad \min_{\boldsymbol{x}\in\mathcal{R}} \quad \max_{\boldsymbol{u}\in\mathcal{U}(\boldsymbol{\iota},\boldsymbol{s})} \quad \left\{ \max_{\boldsymbol{x}'\in\mathcal{R}} \; \boldsymbol{u}^\top\boldsymbol{x}' - \boldsymbol{u}^\top\boldsymbol{x} \right\}. \tag{$\mathcal{P}^K_{\text{off,regret}}$}$$

The following lemma shows that, in a manner paralleling the risk averse case, Problem ($\mathcal{P}^K_{\text{off,regret}}$) can be considerably simplified by eliminating the dependence of $\mathcal{S}(\boldsymbol{\iota})$ on $\boldsymbol{\iota}$, by dropping the "indifferent" scenarios, and by replacing the strict inequalities in the set $\mathcal{U}(\boldsymbol{\iota},\boldsymbol{s})$ by their loose counterparts.

Lemma 4. *Problem ($\mathcal{P}^K_{\text{off,regret}}$) is equivalent to*

$$\underset{\boldsymbol{\iota}\in\mathcal{C}^K}{\text{minimize}} \quad \max_{\boldsymbol{s}\in\widetilde{\mathcal{S}}^K} \quad \min_{\boldsymbol{x}\in\mathcal{R}} \quad \max_{\boldsymbol{u}\in\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s})} \quad \left\{ \max_{\boldsymbol{x}'\in\mathcal{R}} \; \boldsymbol{u}^\top\boldsymbol{x}' - \boldsymbol{u}^\top\boldsymbol{x} \right\}, \tag{$\widetilde{\mathcal{P}}^K_{\text{off,regret}}$}$$

in the sense that the two problems have the same optimal objective value and the same sets of optimal solutions.

Next, we show that Problem ($\widetilde{\mathcal{P}}^K_{\text{off,regret}}$) is generally $\mathcal{NP}$-hard when the recommendation set is discrete, which motivates the integer programming based reformulations we provide in the next sections.

Theorem 5. *Evaluating the objective function of Problem ($\widetilde{\mathcal{P}}^K_{\text{off,regret}}$) is $\mathcal{NP}$-hard even if the recommendation set $\mathcal{R}$ consists of only two elements.*

Motivated by the complexity results above, we provide an enumeration-based approach for solving ($\widetilde{\mathcal{P}}^K_{\text{off,regret}}$) in Section 5.2 applicable to cases where K is small. We provide a general, MBLP reformulation of Problem ($\widetilde{\mathcal{P}}^K_{\text{off,regret}}$) in Section 5.3 and a column-and-constraint generation approach in Section 5.4. These are motivated by settings where K is larger (in the order of 8 to 10). Throughout the remainder of this section, we make the assumption that the set $\mathcal{R}$ has fixed finite cardinality.

5.2. An Enumeration-Based Solution Approach for Small K

In the case of a small number of queries K, our approach relies on the following proposition.

Algorithm 3: Enumeration-based algorithm for solving $(\widetilde{\mathcal{P}}^K_{\text{off,regret}})$ for moderate values of K.

Input: Comparison set $\mathcal{C}$, recommendation set $\mathcal{R}$, uncertainty set $\mathcal{U}^0$, number of queries K;

Output: Optimal query $\boldsymbol{\iota}^\star$ from $\mathcal{C}^K$;

Initialization: $\boldsymbol{\iota}^\star \leftarrow \emptyset$; $\text{OPT} \leftarrow +\infty$;

for $\boldsymbol{\iota} \in \mathcal{C}^K$ **do**

  foreach $\boldsymbol{s} \in \widetilde{\mathcal{S}}^K$ **do**

    $\text{OPT}(\boldsymbol{\iota}, \boldsymbol{s}) \leftarrow$ optimal objective of Problem (7);

  end

  if $\max_{\boldsymbol{s} \in \widetilde{\mathcal{S}}^K} \text{OPT}(\boldsymbol{\iota}, \boldsymbol{s}) < \text{OPT}$ **then**

    $\text{OPT} \leftarrow \max_{\boldsymbol{s} \in \widetilde{\mathcal{S}}^K} \text{OPT}(\boldsymbol{\iota}, \boldsymbol{s})$; $\boldsymbol{\iota}^\star \leftarrow \boldsymbol{\iota}$;

  end

end

Proposition 3. *For any fixed* $\boldsymbol{\iota} \in \mathcal{C}^K$ *and* $\boldsymbol{s} \in \widetilde{\mathcal{S}}^K$*, the regret averse recommendation problem* $(\mathcal{R}_{\text{regret}})$ *is equivalent to the minimization problem*

$$
\begin{array}{lll}
\text{minimize} & \theta & \\
\text{subject to} & \theta \in \mathbb{R},\ \boldsymbol{x} \in \mathcal{R} & \\
& \left.\begin{array}{l}
\boldsymbol{\alpha}^{\boldsymbol{x}'} \in \mathbb{R}^K_-,\ \boldsymbol{\beta}^{\boldsymbol{x}'} \in \mathbb{R}^M_- \\
\theta \geq \boldsymbol{b}^\top \boldsymbol{\beta}^{\boldsymbol{x}'} \\
\displaystyle\sum_{\kappa \in \mathcal{K}} \boldsymbol{s}_\kappa \left(\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}\right) \boldsymbol{\alpha}_\kappa^{\boldsymbol{x}'} + \boldsymbol{B}^\top \boldsymbol{\beta}^{\boldsymbol{x}'} = \boldsymbol{x}' - \boldsymbol{x}
\end{array}\right\} & \forall \boldsymbol{x}' \in \mathcal{R},
\end{array}
\tag{7}
$$

whose size is polynomial in the size of the input. Problem (7) *is an MBLP if* $\mathcal{R}$ *has fixed finite cardinality.*

Similar to the max-min utility case, if K is small, we leverage Proposition 3 and the moderate cardinality of $\mathcal{C}^K$ and $\widetilde{\mathcal{S}}^K$ to devise an enumeration-based algorithm for solving Problem $(\widetilde{\mathcal{P}}^K_{\text{off,regret}})$. For each $\boldsymbol{\iota} \in \mathcal{C}^K$, we solve 2^K instances of Problem (7) (one for each element $\boldsymbol{s}$ of $\widetilde{\mathcal{S}}^K$) and record the maximum value of the objective over these instances. The optimal solution to Problem $(\widetilde{\mathcal{P}}^K_{\text{off,regret}})$ is given by the choice of $\boldsymbol{\iota}$ that gives the smallest maximum value. The total number of problems solved is $(2|\mathcal{C}|)^K$ and enumeration can thus be done efficiently if K is small (and I is not too large). The steps to follow are summarized in Algorithm 3.

5.3. Exact MBLP Reformulation

As in the max-min utility case, the enumeration approach from Section 5.2 will not scale to realistic values of K (8 to 10). In this section, we thus propose an exact MBLP reformulation for Problem ($\widetilde{\mathcal{P}}^K_{\rm off,regret}$), which enables us to select queries optimally even when K is large by leveraging techniques from integer optimization to speed-up computation and circumvent complete enumeration (see also Section 7).

Lemma 5. *Problem ($\widetilde{\mathcal{P}}^K_{\rm off,regret}$) is equivalent to the following finite program*

$$
\begin{array}{ll}
\text{minimize} & \tau \\
\text{subject to} & \tau \in \mathbb{R},\ \boldsymbol{\iota} \in \mathcal{C}^K,\ \boldsymbol{x}^{\boldsymbol{s}} \in \mathcal{R},\ \boldsymbol{s} \in \widetilde{\mathcal{S}}^K \\
& \left.\begin{array}{l}
\boldsymbol{\alpha}^{(\boldsymbol{x}',\boldsymbol{s})} \in \mathbb{R}^K_-,\ \boldsymbol{\beta}^{(\boldsymbol{x}',\boldsymbol{s})} \in \mathbb{R}^M_- \\
\tau \ \geq \ \boldsymbol{b}^\top \boldsymbol{\beta}^{(\boldsymbol{x}',\boldsymbol{s})} \\
\displaystyle\sum_{\kappa \in \mathcal{K}} \boldsymbol{s}_\kappa \left(\boldsymbol{x}^{\iota_1^\kappa} - \boldsymbol{x}^{\iota_2^\kappa}\right) \boldsymbol{\alpha}^{(\boldsymbol{x}',\boldsymbol{s})}_\kappa + \boldsymbol{B}^\top \boldsymbol{\beta}^{(\boldsymbol{x}',\boldsymbol{s})} = \boldsymbol{x}' - \boldsymbol{x}^{\boldsymbol{s}}
\end{array}\right\} \quad \forall \boldsymbol{s} \in \widetilde{\mathcal{S}}^K,\ \forall \boldsymbol{x}' \in \mathcal{R},
\end{array}
\tag{8}
$$

where $\boldsymbol{\iota}$ denotes the queries to make and $\boldsymbol{x}^{\boldsymbol{s}}$ the items to recommend in response scenario $\boldsymbol{s} \in \widetilde{\mathcal{S}}^K$.

Next, we convert Problem (8) to a finite mixed-binary linear program by following a similar approach to that described in Section 3.3. The result is summarized in the following theorem.

Theorem 6. *Problem* (8) *is equivalent to the following mixed-binary linear program*

$$
\begin{array}{ll}
\text{minimize} & \tau \\
\text{subject to} & \tau \in \mathbb{R},\ \boldsymbol{v}^\kappa,\ \boldsymbol{w}^\kappa \in \{0,1\}^I,\ \kappa \in \mathcal{K} \\
& \boldsymbol{x}^{\boldsymbol{s}} \in \mathcal{R} \quad \forall \boldsymbol{s} \in \widetilde{\mathcal{S}}^K \\
& \boldsymbol{\alpha}^{(\boldsymbol{x}',\boldsymbol{s})} \in \mathbb{R}^K_-,\ \boldsymbol{\beta}^{(\boldsymbol{x}',\boldsymbol{s})} \in \mathbb{R}^M_- \quad \forall \boldsymbol{s} \in \widetilde{\mathcal{S}}^K,\ \boldsymbol{x}' \in \mathcal{R} \\
& \overline{\boldsymbol{v}}^{(\boldsymbol{x}',\boldsymbol{s})\kappa},\ \overline{\boldsymbol{w}}^{(\boldsymbol{x}',\boldsymbol{s})\kappa} \in \mathbb{R}^I_- \quad \forall \boldsymbol{s} \in \widetilde{\mathcal{S}}^K,\ \boldsymbol{x}' \in \mathcal{R},\ \kappa \in \mathcal{K} \\
& \left.\begin{array}{l}
\tau \ \geq \ \boldsymbol{b}^\top \boldsymbol{\beta}^{(\boldsymbol{x}',\boldsymbol{s})} \\
\displaystyle\sum_{i \in \mathcal{I}} \boldsymbol{x}^i \sum_{\kappa \in \mathcal{K}} \boldsymbol{s}_\kappa \left(\overline{\boldsymbol{v}}_i^{(\boldsymbol{x}',\boldsymbol{s})\kappa} - \overline{\boldsymbol{w}}_i^{(\boldsymbol{x}',\boldsymbol{s})\kappa}\right) + \boldsymbol{B}^\top \boldsymbol{\beta}^{(\boldsymbol{x}',\boldsymbol{s})} \ = \ \boldsymbol{x}' - \boldsymbol{x}^{\boldsymbol{s}}
\end{array}\right\} \quad \begin{array}{c} \forall \boldsymbol{s} \in \widetilde{\mathcal{S}}^K, \\ \boldsymbol{x}' \in \mathcal{R} \end{array} \\
& \mathbf{e}^\top \boldsymbol{v}^\kappa = 1,\ \mathbf{e}^\top \boldsymbol{w}^\kappa = 1,\ 1 - \boldsymbol{w}^\kappa_i \ \geq \displaystyle\sum_{i' : i' \geq i} \boldsymbol{v}^\kappa_{i'} \quad \forall i \in \mathcal{I},\ \forall \kappa \in \mathcal{K} \\
& \left.\begin{array}{l}
\overline{\boldsymbol{v}}^{(\boldsymbol{x}',\boldsymbol{s})\kappa} \ \geq \ -M \boldsymbol{v}^\kappa,\ \overline{\boldsymbol{v}}^{(\boldsymbol{x}',\boldsymbol{s})\kappa} \ \geq \ \boldsymbol{\alpha}^{(\boldsymbol{x}',\boldsymbol{s})}_\kappa \mathbf{e} \\
\overline{\boldsymbol{v}}^{(\boldsymbol{x}',\boldsymbol{s})\kappa} \ \leq \ \boldsymbol{\alpha}^{(\boldsymbol{x}',\boldsymbol{s})}_\kappa \mathbf{e} + M(\mathbf{e} - \boldsymbol{v}^\kappa) \\
\overline{\boldsymbol{w}}^{(\boldsymbol{x}',\boldsymbol{s})\kappa} \ \geq \ -M \boldsymbol{w}^\kappa,\ \overline{\boldsymbol{w}}^{(\boldsymbol{x}',\boldsymbol{s})\kappa} \ \geq \ \boldsymbol{\alpha}^{(\boldsymbol{x}',\boldsymbol{s})}_\kappa \mathbf{e} \\
\overline{\boldsymbol{w}}^{(\boldsymbol{x}',\boldsymbol{s})\kappa} \ \leq \ \boldsymbol{\alpha}^{(\boldsymbol{x}',\boldsymbol{s})}_\kappa \mathbf{e} + M(\mathbf{e} - \boldsymbol{w}^\kappa)
\end{array}\right\} \quad \begin{array}{c} \forall \boldsymbol{s} \in \widetilde{\mathcal{S}}^K,\ \boldsymbol{x}' \in \mathcal{R}, \\ \kappa \in \mathcal{K}, \end{array}
\end{array}
\tag{9}
$$

where M is a "big-M" constant. In particular, given an optimal solution $(\tau, \boldsymbol{v}, \boldsymbol{w}, \boldsymbol{\alpha}, \boldsymbol{\beta}, \overline{\boldsymbol{v}}, \overline{\boldsymbol{w}})$ to Problem (9), an optimal set of queries for the risk averse offline preference elicitation problem is given by

$$\boldsymbol{\iota}_1^\kappa = \sum_{i \in \mathcal{I}} i \cdot \mathbb{I}\left(\boldsymbol{v}_i^\kappa = 1\right) \quad \textit{and} \quad \boldsymbol{\iota}_2^\kappa = \sum_{i \in \mathcal{I}} i \cdot \mathbb{I}\left(\boldsymbol{w}_i^\kappa = 1\right), \quad \kappa \in \mathcal{K}.$$

The first and second sets of constraints in Problem (9) originate from the first two constraints in (8). The third set of constraints ensure that $\boldsymbol{\iota}^\kappa \in \mathcal{C}$ for all κ. The remaining constraints are used to linearize the products $\boldsymbol{v}^\kappa \boldsymbol{\alpha}_\kappa^{(\boldsymbol{x}',\boldsymbol{s})}$ and $\boldsymbol{w}^\kappa \boldsymbol{\alpha}_\kappa^{(\boldsymbol{x}',\boldsymbol{s})}$. Reformulation (9) is very attractive as it enables us to solve the regret averse offline preference elicitation problem ($\mathcal{P}^K_{\text{off,regret}}$) as an MBLP using off-the-shelf solvers. For fixed K, Problem (9) is polynomial in the size of the input. Yet, it is exponential in K. In the next section, we propose a column-and-constraint generation approach, similar to the one proposed in Section 3.4 for the max-min utility case, that enables us to scale to practical values of K.

5.4. Column-and-Constraint Generation

Problem (9) presents a number of decision variables and constraints that are exponential in K. Similar to the max-min utility case (see Section 3.4), we leverage the robust nature of the problem and propose a column-and-constraint generation algorithm to speed-up computation. We describe our column-and-constraint generation procedure using the finite program (8). Naturally, all problems solved would need to be converted to MBLPs first, using techniques similar to those employed in Theorem 6. We omit these conversions to streamline presentation.

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

$$\begin{array}{ll}
\text{minimize} & \tau \\
\text{subject to} & \tau \in \mathbb{R},\ \boldsymbol{\iota} \in \mathcal{C}^K,\ \boldsymbol{x}^{\boldsymbol{s}} \in \mathcal{R},\ \boldsymbol{s} \in \widetilde{\mathcal{S}}^K \\
& \left.\begin{array}{l}
\boldsymbol{\alpha}^{(\boldsymbol{x}',\boldsymbol{s})} \in \mathbb{R}^K_-,\ \boldsymbol{\beta}^{(\boldsymbol{x}',\boldsymbol{s})} \in \mathbb{R}^M_- \\
\tau \ \geq \ \boldsymbol{b}^\top \boldsymbol{\beta}^{(\boldsymbol{x}',\boldsymbol{s})} \\
\displaystyle\sum_{\kappa \in \mathcal{K}} \boldsymbol{s}_\kappa \left(\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}\right) \boldsymbol{\alpha}_\kappa^{(\boldsymbol{x}',\boldsymbol{s})} + \boldsymbol{B}^\top \boldsymbol{\beta}^{(\boldsymbol{x}',\boldsymbol{s})} \ = \ \boldsymbol{x}' - \boldsymbol{x}^{\boldsymbol{s}}
\end{array}\right\} \begin{array}{l} \forall \boldsymbol{s} \in \mathcal{S}', \\ \boldsymbol{x}' \in \mathcal{R}. \end{array}
\end{array} \qquad (\mathcal{CCG}^{\text{master}}_{\text{regret}}(\mathcal{S}'))$$

This problem only involves a subset of the decision variables and constraints of Problem (8) (those indexed by $\boldsymbol{s} \in \mathcal{S}' \subseteq \widetilde{\mathcal{S}}^K$). Given variables $(\tau, \boldsymbol{\iota})$ feasible in the master problem, we define the $(\boldsymbol{x}', \boldsymbol{s})$th subproblem, $\boldsymbol{x}' \in \mathcal{R}$, $\boldsymbol{s} \in \widetilde{\mathcal{S}}^K$, through

$$\begin{array}{ll} \text{minimize} & 0 \\ \text{subject to} & \boldsymbol{\alpha} \in \mathbb{R}^K_-,\ \boldsymbol{\beta} \in \mathbb{R}^M_-,\ \boldsymbol{x} \in \mathcal{R} \\ & \tau \ \geq \ \boldsymbol{b}^\top \boldsymbol{\beta} \\ & \displaystyle\sum_{\kappa \in \mathcal{K}} \boldsymbol{s}_\kappa \left(\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}\right) \boldsymbol{\alpha}_\kappa + \boldsymbol{B}^\top \boldsymbol{\beta} \ = \ \boldsymbol{x}' - \boldsymbol{x}. \end{array} \qquad (\mathcal{CCG}^{\text{sub},(\boldsymbol{x}',\boldsymbol{s})}_{\text{regret}}(\tau, \boldsymbol{\iota}))$$

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

$$\begin{array}{ll} \text{maximize} & \theta \\ \text{subject to} & \theta \in \mathbb{R},\ \boldsymbol{s} \in \widetilde{\mathcal{S}}^K \\ & \boldsymbol{x}'^{,\boldsymbol{x}} \in \mathcal{R},\ \boldsymbol{u}^{\boldsymbol{x}} \in \mathbb{R}^J \quad \forall \boldsymbol{x} \in \mathcal{R} \\ & \theta \ \leq \ (\boldsymbol{u}^{\boldsymbol{x}})^\top (\boldsymbol{x}'^{,\boldsymbol{x}} - \boldsymbol{x}) \quad \forall \boldsymbol{x} \in \mathcal{R} \\ & \left. \begin{array}{l} (\boldsymbol{u}^{\boldsymbol{x}})^\top (\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}) \ \geq \ -M(1 - \boldsymbol{s}_\kappa) \\ (\boldsymbol{u}^{\boldsymbol{x}})^\top (\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}) \ \leq \ M(\boldsymbol{s}_\kappa + 1) \end{array} \right\} \quad \forall \boldsymbol{x} \in \mathcal{R}. \end{array} \qquad (\mathcal{CCG}^{\text{feas}}_{\text{regret}}(\boldsymbol{\iota}))$$

Proposition 4. *Let $\boldsymbol{\iota}$ be feasible in the relaxed master problem $(\mathcal{CCG}^{\text{master}}_{\text{regret}}(\mathcal{S}'))$. Then, $\boldsymbol{\iota}$ is feasible in Problem $(\widetilde{\mathcal{P}}^K_{\text{off,regret}})$ and the objective value of $\boldsymbol{\iota}$ in Problem $(\widetilde{\mathcal{P}}^K_{\text{off,regret}})$ is given by the optimal objective value of Problem $(\mathcal{CCG}^{\text{feas}}_{\text{regret}}(\boldsymbol{\iota}))$. Moreover, if the set $\mathcal{R}$ has fixed finite cardinality, then Problem $(\mathcal{CCG}^{\text{feas}}_{\text{regret}}(\boldsymbol{\iota}))$ is a mixed-binary linear program of size polynomial in the size of the input.*

Lemma 6. *The relaxed master problem $(\mathcal{CCG}^{\text{master}}_{\text{regret}}(\mathcal{S}'))$ is always feasible. If $(\mathcal{CCG}^{\text{master}}_{\text{regret}}(\mathcal{S}'))$ is solvable, let $(\tau, \boldsymbol{\iota}, \{\boldsymbol{\alpha}^{(\boldsymbol{x}',\boldsymbol{s})}, \boldsymbol{\beta}^{(\boldsymbol{x}',\boldsymbol{s})}\}_{\boldsymbol{x}' \in \mathcal{R}, \boldsymbol{s} \in \mathcal{S}'}, \{\boldsymbol{x}^{\boldsymbol{s}}\}_{\boldsymbol{s} \in \mathcal{S}'})$ be an optimal solution. Else, if $(\mathcal{CCG}^{\text{master}}_{\text{regret}}(\mathcal{S}'))$ is unbounded, set $\tau = -\infty$ and let $\boldsymbol{\iota} \in \mathcal{C}^K$ be such that $(\mathcal{CCG}^{\text{master}}_{\text{regret}}(\mathcal{S}'))$ is unbounded when $\boldsymbol{\iota}$ is fixed to that value. Moreover, let $(\theta, \{\boldsymbol{u}^{\boldsymbol{x}}, \boldsymbol{x}'^{,\boldsymbol{x}}\}_{\boldsymbol{x} \in \mathcal{R}}, \boldsymbol{s})$ be optimal in Problem $(\mathcal{CCG}^{\text{feas}}_{\text{regret}}(\boldsymbol{\iota}))$. Then, the following hold:*

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

(ii) *If $\theta = \tau$, then Problem $(\mathcal{CCG}^{\text{sub},(\boldsymbol{x}',\boldsymbol{s})}_{\text{regret}}(\tau, \boldsymbol{\iota}))$ is feasible for all $\boldsymbol{x}' \in \mathcal{R}$ and $\boldsymbol{s} \in \widetilde{\mathcal{S}}^K$;*

(iii) *If $\theta > \tau$, then there exists $\boldsymbol{x}' \in \mathcal{R}$ such that the pair $(\boldsymbol{x}', \boldsymbol{s})$ corresponds to an infeasible subproblem, i.e., Problem $(\mathcal{CCG}^{\text{sub},(\boldsymbol{x}',\boldsymbol{s})}_{\text{regret}}(\tau, \boldsymbol{\iota}))$ is infeasible.*

Algorithm 4: Column-and-Constraint Generation procedure for solving Problem ($\widetilde{\mathcal{P}}^K_{\rm off,regret}$).

Inputs: Optimality tolerance δ, comparison set $\mathcal{C}$, and recommendation set $\mathcal{R}$;

Initial uncertainty set $\mathcal{U}^0$ and number of queries K;

Output: Query $\boldsymbol{\iota}^\star$ from $\mathcal{C}^K$, near optimal in Problem ($\widetilde{\mathcal{P}}^K_{\rm off,regret}$) with associated objective θ;

Initialization: $\boldsymbol{\iota}^\star \leftarrow \emptyset$; Upper and lower bounds: $\text{UB} \leftarrow +\infty$ and $\text{LB} \leftarrow -\infty$;

Initialize index set: $\mathcal{S}' \leftarrow \emptyset$;

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

Solve the master problem ($\mathcal{CCG}^{\rm master}_{\rm regret}(\mathcal{S}')$). If it is solvable, let

($\tau, \boldsymbol{\iota}, \{\boldsymbol{\alpha}^{(\boldsymbol{x}',\boldsymbol{s})}, \boldsymbol{\beta}^{(\boldsymbol{x}',\boldsymbol{s})}\}_{\boldsymbol{x}' \in \mathcal{R}, \boldsymbol{s} \in \mathcal{S}'}, \{\boldsymbol{x}^{\boldsymbol{s}}\}_{\boldsymbol{s} \in \mathcal{S}'}$) be an optimal solution. If it is unbounded, set $\tau = -\infty$

and let $\boldsymbol{\iota} \in \mathcal{C}^K$ be such that ($\mathcal{CCG}^{\rm master}_{\rm regret}(\mathcal{S}')$) is unbounded when $\boldsymbol{\iota}$ is fixed to that value;

Set $\text{LB} \leftarrow \tau$;

Solve feasibility subproblem ($\mathcal{CCG}^{\rm feas}_{\rm regret}(\boldsymbol{\iota})$). Let ($\theta, \{\boldsymbol{u}^{\boldsymbol{x}}, \boldsymbol{x}'^{,\boldsymbol{x}}\}_{\boldsymbol{x} \in \mathcal{R}}, \boldsymbol{s}$) denote an optimal solution;

Set $\text{UB} \leftarrow \theta$;

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

$\mathcal{S}' \leftarrow \mathcal{S}' \cup \{\boldsymbol{s}\}$

end

end

Set $\boldsymbol{\iota}^\star \leftarrow \boldsymbol{\iota}$;

Result: Collection of queries $\boldsymbol{\iota}^\star$ near-optimal in ($\widetilde{\mathcal{P}}^K_{\rm off,regret}$) with objective value θ.

Propositions 4 and Lemma 6 culminate in Algorithm 4 whose convergence is guaranteed by the following theorem.

Theorem 7. *Algorithm 4 terminates in a final number of steps with a feasible solution to the regret averse active preference elicitation Problem ($\widetilde{\mathcal{P}}^K_{\rm off,regret}$). The objective value attained by this solution is within δ of the optimal objective value of the problem.*

5.5. Active Online Elicitation with the Min-Max Regret Decision Criterion

In Sections 5.1-5.4, we assumed that all queries are chosen at once. In many settings of practical interest, queries are selected one at a time and the answer to each query is revealed before the next query is selected.

In this case, the recommender system has the opportunity to *adjust* the queries made, taking into account the answers to past queries. In this section, we address this *regret averse online active preference elicitation problem* (also referred to as *online active preference elicitation problem with the max-min utility decision criterion*) where comparisons are selected adaptively over time.

Mathematically, the online regret averse recommendation problem is expressible as

$$\max_{\boldsymbol{\iota}^1 \in \mathcal{C}} \min_{\boldsymbol{s}_1 \in \mathcal{S}^1(\boldsymbol{\iota}^1)} \cdots \max_{\boldsymbol{\iota}^K \in \mathcal{C}} \min_{\boldsymbol{s}_K \in \mathcal{S}^K(\boldsymbol{\iota}^{[K]}, \boldsymbol{s}_{[K-1]})} \max_{\boldsymbol{x} \in \mathcal{R}} \min_{\boldsymbol{u} \in \mathcal{U}(\boldsymbol{\iota}, \boldsymbol{s})} \left\{ \max_{\boldsymbol{x}' \in \mathcal{R}} \boldsymbol{u}^\top \boldsymbol{x}' - \boldsymbol{u}^\top \boldsymbol{x} \right\}. \qquad (\mathcal{P}^K_{\text{on,regret}})$$

As in the max-min utility case, we propose to approximate the adaptive query decisions by constant decision rules, allowing the recommendation decisions to be fully-adaptive, and solving the problem in a folding horizon fashion, see Section 4.2.

6. Active Preference Elicitation under Inconsistent Responses

In this section, we generalize all the ideas in the paper to cases where the user may be boundedly rational, giving answers that are inconsistent with Assumptions 1 and 2.

6.1. Inconsistencies Model

So far, we have assumed that users are perfectly rational and maximize the expectation of a utility in the spirit of expected utility theory, as claimed by Neumann and Morgenstern (1944), see Assumption 1. Yet, several authors have disputed this claim and shown that oftentimes, individuals behave in seemingly "irrational" ways, see e.g., Kahneman and Tversky (1979), Kahneman et al. (1991), Allais (1953). In particular, these works 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 addition, we have assumed that the user utility function is linear with known attributes, see Assumption 2. While this assumption is very common in the literature, it may not necessarily hold true in practice. Therefore, in this section, we relax Assumptions 1 and 2. Noisy/inconsistent responses have been investigated in the literature in the context of polyhedral methods to preference elicitation, see e.g., Toubia et al. (2003) and Bertsimas and O'Hair (2013). Our approach differs from these works in that we are able to integrate the learning and the downstream optimization, while still taking into account the possibility of inconsistencies in the user responses.

We propose to capture the possibility that user responses are inconsistent or incompatible with the assumptions made by interpreting user responses with some "reservation". We take the viewpoint that, when comparing (the utility of) items ι_1^κ and ι_2^κ, $\kappa \in \mathcal{K}$, user answers are perturbed by additive noise. Thus, rather than being based on the sign of the difference $(\mathscr{U}^\top \boldsymbol{x}^{\iota_1^\kappa} - \mathscr{U}^\top \boldsymbol{x}^{\iota_2^\kappa})$, their answer is based on the sign of $(\mathscr{U}^\top \boldsymbol{x}^{\iota_1^\kappa} - \mathscr{U}^\top \boldsymbol{x}^{\iota_2^\kappa} + \mathscr{E}_\kappa)$, where $\mathscr{E} \in \mathbb{R}^K$ is a vector of independent identically distributed errors. This idea is similar in spirit to that originally proposed by Toubia et al. (2003).

In line with modern robust optimization, see e.g., Ben-Tal et al. (2009), Bertsimas et al. (2010), Gorissen et al. (2015), we model the random parameters $\mathscr{E}$ as decision variables $\boldsymbol{\epsilon}$ constrained to lie in an uncertainty set, which we denote by $\mathcal{E}_\Gamma$. Under this model, we relax Assumptions 1 and 2 as follows.

Assumption 4 (Noisy Preferences). *We assume that the user preferences over items satisfy:*

1. *If* $\boldsymbol{x} \succeq \boldsymbol{y}$, *then* $\mathscr{U}(\boldsymbol{x}) - \mathscr{U}(\boldsymbol{y}) + \mathscr{E}_\kappa \geq 0$; *and*
2. *If* $\boldsymbol{x} \sim \boldsymbol{y}$, *then* $|\mathscr{U}(\boldsymbol{x}) - \mathscr{U}(\boldsymbol{y})| \leq \mathscr{E}_\kappa$.

Moreover, the user assumed utility function $u : \mathcal{X} \to \mathbb{R}$ *is expressible as* $u(\boldsymbol{x}) := \mathscr{U}^\top \boldsymbol{x}$ *for some (random) vector* $\mathscr{U}$ *supported in the non-empty and bounded uncertainty set* $\mathcal{U}^0 := \{\boldsymbol{u} \in \mathbb{R}^J \mid \boldsymbol{B}\boldsymbol{u} \geq \boldsymbol{b}\}$ *for some (known) matrix* $\boldsymbol{B} \in \mathbb{R}^{M \times J}$ *and vector* $\boldsymbol{b} \in \mathbb{R}^M$. *The user inconsistencies across* K *queries are assumed to lie in the uncertainty set* $\mathcal{E}_\Gamma := \{\boldsymbol{\epsilon} \in \mathbb{R}_+^K : \sum_{\kappa \in \mathcal{K}} \boldsymbol{\epsilon}_\kappa \leq \Gamma\}$ *for some (known) parameter* $\Gamma \in \mathbb{R}_+$.

Note that Assumption 4 directly generalizes Assumptions 1 and 2. Indeed, when $\Gamma = 0$, Assumption 4 reduces to Assumptions 1 and 2. Uncertainty sets of the type $\mathcal{E}_\Gamma$ are very popular in the literature, see e.g., Bandi and Bertsimas (2012). The assumption on $\mathcal{E}_\Gamma$ can easily be relaxed to merely requiring that $\mathcal{E}_\Gamma$ be a bounded full-dimensional polyhedron.

6.2. Risk Averse & Regret Averse Recommendations under Inconsistent Responses

Given the sequences $\boldsymbol{\iota} \in \mathcal{C}^K$ and $\boldsymbol{s} \in \widetilde{\mathcal{S}}^K$ of questions and answers, the recommender system offers the item with the maximum worst-case (minimum) utility for any $\boldsymbol{u} \in \mathcal{U}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s})$. Under Assumption 4, the risk averse recommendation problem reads

$$\underset{\boldsymbol{x} \in \mathcal{R}}{\text{maximize}} \quad \min_{\boldsymbol{u} \in \mathcal{U}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s})} \quad \boldsymbol{u}^\top \boldsymbol{x}, \tag{$\mathcal{R}_{\text{risk}}^\Gamma$}$$

where,

$$\mathcal{U}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s}) := \left\{ \begin{array}{ll} \boldsymbol{u} \in \mathcal{U}^0 : \exists \boldsymbol{\epsilon} \in \mathcal{E}_\Gamma \text{ such that} & \\ \qquad \boldsymbol{u}^\top (\boldsymbol{x}^{\iota_1^\kappa} - \boldsymbol{x}^{\iota_2^\kappa}) \geq -\boldsymbol{\epsilon}_\kappa & \forall \kappa \in \mathcal{K} : \boldsymbol{s}_\kappa = 1 \\ \qquad |\boldsymbol{u}^\top (\boldsymbol{x}^{\iota_1^\kappa} - \boldsymbol{x}^{\iota_2^\kappa})| \leq \boldsymbol{\epsilon}_\kappa & \forall \kappa \in \mathcal{K} : \boldsymbol{s}_\kappa = 0 \\ \qquad \boldsymbol{u}^\top (\boldsymbol{x}^{\iota_1^\kappa} - \boldsymbol{x}^{\iota_2^\kappa}) \leq \boldsymbol{\epsilon}_\kappa & \forall \kappa \in \mathcal{K} : \boldsymbol{s}_\kappa = -1 \end{array} \right\}.$$

Accordingly, the regret averse recommendation problem is expressible as

$$\underset{\boldsymbol{x}\in\mathcal{R}}{\text{minimize}} \quad \max_{\boldsymbol{u}\in\mathcal{U}_\Gamma(\boldsymbol{\iota},\boldsymbol{s})} \quad \left\{ \max_{\boldsymbol{x}'\in\mathcal{R}} \; \boldsymbol{u}^\top\boldsymbol{x}' - \boldsymbol{u}^\top\boldsymbol{x} \right\}. \tag{$\mathcal{R}^\Gamma_{\text{regret}}$}$$

6.3. Offline Risk Averse Active Preference Elicitation under Inconsistencies

The offline risk averse recommendation problem with inconsistencies is expressible as the following two-stage robust optimization problem with decision-dependent information discovery

$$\underset{\boldsymbol{\iota}\in\mathcal{C}^K}{\text{maximize}} \quad \min_{\boldsymbol{s}\in\mathcal{S}_\Gamma(\boldsymbol{\iota})} \quad \max_{\boldsymbol{x}\in\mathcal{R}} \quad \min_{\boldsymbol{u}\in\mathcal{U}_\Gamma(\boldsymbol{\iota},\boldsymbol{s})} \quad \boldsymbol{u}^\top\boldsymbol{x}, \tag{$\mathcal{P}^{\Gamma,K}_{\text{off,risk}}$}$$

where $\mathcal{S}_\Gamma(\boldsymbol{\iota}) := \{\boldsymbol{s}\in\mathcal{S}^K \mid \mathcal{U}_\Gamma(\boldsymbol{\iota},\boldsymbol{s})\neq\emptyset\}$ denotes the set of all answers compatible with some $\boldsymbol{u}\in\mathcal{U}^0$ and $\boldsymbol{\epsilon}\in\mathcal{E}_\Gamma$. As before, we can simplify the problem above by eliminating the dependence of $\mathcal{S}_\Gamma(\boldsymbol{\iota})$ on $\boldsymbol{\iota}$ and by noting that "indifferent" answers can always be omitted. In the case of a moderate number of queries K, we can leverage the following proposition to solve the risk averse active preference elicitation problem with inconsistencies ($\mathcal{P}^{\Gamma,K}_{\text{off,risk}}$) by enumeration, in a way that parallels Algorithm 1.

Proposition 5. *For any fixed $\boldsymbol{\iota}\in\mathcal{C}^K$ and $\boldsymbol{s}\in\widetilde{\mathcal{S}}^K$, the robust recommendation problem ($\mathcal{R}^\Gamma_{\text{risk}}$) is equivalent to the maximization problem*

$$\begin{aligned} \text{maximize} \quad & \boldsymbol{b}^\top\boldsymbol{\beta} + \Gamma\mu \\ \text{subject to} \quad & \boldsymbol{x}\in\mathcal{R},\; \boldsymbol{\alpha}\in\mathbb{R}^K_+,\; \boldsymbol{\beta}\in\mathbb{R}^M_+,\; \mu\in\mathbb{R}_- \\ & \sum_{\kappa\in\mathcal{K}} \boldsymbol{s}_\kappa \left(\boldsymbol{x}^{\boldsymbol{\iota}^\kappa_1} - \boldsymbol{x}^{\boldsymbol{\iota}^\kappa_2}\right)\boldsymbol{\alpha}_\kappa + \boldsymbol{B}^\top\boldsymbol{\beta} \;=\; \boldsymbol{x} \\ & \boldsymbol{\alpha} + \mu\mathbf{e} \;\leq\; \mathbf{0}, \end{aligned} \tag{10}$$

whose size is polynomial in the size of the input. Problem (10) *is a linear program if $\mathcal{R}$ is polyhedral and a mixed-binary linear program if $\mathcal{R}$ also involves integrality constraints.*

When K is larger than 2 or 3, solving Problem ($\mathcal{P}^{\Gamma,K}_{\text{off,risk}}$) by enumeration becomes computationally prohibitive. In such settings, the following lemma and theorem can be leveraged to solve ($\mathcal{P}^{\Gamma,K}_{\text{off,risk}}$) as an MBLP.

Lemma 7. *Problem* $(\mathcal{P}^{\Gamma,K}_{\text{off,risk}})$ *is equivalent to the following finite program*

$$
\begin{array}{ll}
\text{maximize} & \tau \\
\text{subject to} & \tau \in \mathbb{R},\ \boldsymbol{\iota} \in \mathcal{C}^K \\
& \left.\begin{array}{l}
\boldsymbol{x}^{\boldsymbol{s}} \in \mathcal{R},\ \boldsymbol{\alpha}^{\boldsymbol{s}} \in \mathbb{R}^K_+,\ \boldsymbol{\beta}^{\boldsymbol{s}} \in \mathbb{R}^M_+,\ \mu^{\boldsymbol{s}} \in \mathbb{R}_- \\
\tau \le \boldsymbol{b}^\top \boldsymbol{\beta}^{\boldsymbol{s}} + \Gamma \mu^{\boldsymbol{s}} \\
\sum_{\kappa \in \mathcal{K}} \boldsymbol{s}_\kappa \left(\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}\right) \boldsymbol{\alpha}^{\boldsymbol{s}}_\kappa + \boldsymbol{B}^\top \boldsymbol{\beta} = \boldsymbol{x} \\
\boldsymbol{\alpha}^{\boldsymbol{s}} + \mu^{\boldsymbol{s}} \mathbf{e} \le \mathbf{0}
\end{array}\right\} \quad \forall \boldsymbol{s} \in \widetilde{\mathcal{S}}^K,
\end{array}
\tag{11}
$$

where $\boldsymbol{\iota}$ *denotes the queries to make and* $\boldsymbol{x}^{\boldsymbol{s}}$ *the items to recommend in response scenario* $\boldsymbol{s} \in \widetilde{\mathcal{S}}^K$.

Theorem 8. *Problem* (11) *is equivalent to the following mixed-binary linear program*

$$
\begin{array}{ll}
\text{maximize} & \tau \\
\text{subject to} & \tau \in \mathbb{R},\ \boldsymbol{v}^\kappa,\ \boldsymbol{w}^\kappa \in \{0,1\}^I,\ \kappa \in \mathcal{K} \\
& \boldsymbol{\alpha}^{\boldsymbol{s}} \in \mathbb{R}^K_+,\ \boldsymbol{\beta}^{\boldsymbol{s}} \in \mathbb{R}^M_+,\ \mu^{\boldsymbol{s}} \in \mathbb{R}_-,\ \boldsymbol{x}^{\boldsymbol{s}} \in \mathcal{R},\ \boldsymbol{s} \in \widetilde{\mathcal{S}}^K \\
& \overline{\boldsymbol{v}}^{\boldsymbol{s}\kappa},\ \overline{\boldsymbol{w}}^{\boldsymbol{s}\kappa} \in \mathbb{R}^I_+,\ \boldsymbol{s} \in \widetilde{\mathcal{S}}^K,\ \kappa \in \mathcal{K} \\
& \left.\begin{array}{l}
\tau \le \boldsymbol{b}^\top \boldsymbol{\beta}^{\boldsymbol{s}} + \Gamma \mu^{\boldsymbol{s}} \\
\sum_{i \in \mathcal{I}} \boldsymbol{x}^i \sum_{\kappa \in \mathcal{K}} \boldsymbol{s}_\kappa \left(\overline{\boldsymbol{v}}^{\boldsymbol{s}\kappa}_i - \overline{\boldsymbol{w}}^{\boldsymbol{s}\kappa}_i\right) + \boldsymbol{B}^\top \boldsymbol{\beta}^{\boldsymbol{s}} = \boldsymbol{x}^{\boldsymbol{s}} \\
\boldsymbol{\alpha}^{\boldsymbol{s}} + \mu^{\boldsymbol{s}} \mathbf{e} \le \mathbf{0}
\end{array}\right\} \quad \forall \boldsymbol{s} \in \widetilde{\mathcal{S}}^K \\
& \mathbf{e}^\top \boldsymbol{v}^\kappa = 1,\ \mathbf{e}^\top \boldsymbol{w}^\kappa = 1,\ 1 - \boldsymbol{w}^\kappa_i \ge \sum_{i' : i' \ge i} \boldsymbol{v}^\kappa_{i'} \quad \forall i \in \mathcal{I} \quad \forall \kappa \in \mathcal{K} \\
& \left.\begin{array}{l}
\overline{\boldsymbol{v}}^{\boldsymbol{s}\kappa} \le M \boldsymbol{v}^\kappa,\ \overline{\boldsymbol{v}}^{\boldsymbol{s}\kappa} \le \boldsymbol{\alpha}^{\boldsymbol{s}}_\kappa \mathbf{e},\ \overline{\boldsymbol{v}}^{\boldsymbol{s}\kappa} \ge \boldsymbol{\alpha}^{\boldsymbol{s}}_\kappa \mathbf{e} - M(\mathbf{e} - \boldsymbol{v}^\kappa) \\
\overline{\boldsymbol{w}}^{\boldsymbol{s}\kappa} \le M \boldsymbol{w}^\kappa,\ \overline{\boldsymbol{w}}^{\boldsymbol{s}\kappa} \le \boldsymbol{\alpha}^{\boldsymbol{s}}_\kappa \mathbf{e},\ \overline{\boldsymbol{w}}^{\boldsymbol{s}\kappa} \ge \boldsymbol{\alpha}^{\boldsymbol{s}}_\kappa \mathbf{e} - M(\mathbf{e} - \boldsymbol{w}^\kappa)
\end{array}\right\} \quad \begin{array}{l} \forall \boldsymbol{s} \in \widetilde{\mathcal{S}}^K \\ \kappa \in \mathcal{K}, \end{array}
\end{array}
\tag{12}
$$

where M *is a "big-M" constant. In particular, given an optimal solution* $(\tau, \boldsymbol{v}, \boldsymbol{w}, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{x}, \{\mu^{\boldsymbol{s}}\}_{\boldsymbol{s} \in \widetilde{\mathcal{S}}^K}, \overline{\boldsymbol{v}}, \overline{\boldsymbol{w}})$ *to Problem* (12)*, an optimal set of queries for the risk averse offline preference elicitation problem is given by*

$$
\boldsymbol{\iota}^\kappa_1 = \sum_{i \in \mathcal{I}} i \cdot \mathbb{I}\left(\boldsymbol{v}^\kappa_i = 1\right) \quad \text{and} \quad \boldsymbol{\iota}^\kappa_2 = \sum_{i \in \mathcal{I}} i \cdot \mathbb{I}\left(\boldsymbol{w}^\kappa_i = 1\right), \quad \kappa \in \mathcal{K}.
$$

The interpretation of the constraints is similar to the case without inconsistencies discussed in Section 5.3. Problem (12) can be solved using a column-and-constraint generation procedure similar to that proposed in Section 3.4 for the case without inconsistencies. Here, we only provide the master problem, the feasibility subproblem, and the algorithm. To minimize notational overhead, we describe our column-and-constraint generation procedure using the finite program (11). The relaxed master problem is

$$
\begin{array}{lll}
\text{maximize} & \tau & \\
\text{subject to} & \tau \in \mathbb{R},\ \boldsymbol{\iota} \in \mathcal{C}^K & \\
& \left.\begin{array}{l}
\boldsymbol{\alpha}^{\boldsymbol{s}}_{+} \in \mathbb{R}^K,\ \boldsymbol{\beta}^{\boldsymbol{s}} \in \mathbb{R}^M_{+},\ \mu^{\boldsymbol{s}} \in \mathbb{R}_{-},\ \boldsymbol{x}^{\boldsymbol{s}} \in \mathcal{R} \\
\tau \le \boldsymbol{b}^\top \boldsymbol{\beta}^{\boldsymbol{s}} + \Gamma \mu^{\boldsymbol{s}} \\
\displaystyle\sum_{\kappa \in \mathcal{K}} \boldsymbol{s}_\kappa \left(\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}\right) \boldsymbol{\alpha}^{\boldsymbol{s}}_\kappa + \boldsymbol{B}^\top \boldsymbol{\beta}^{\boldsymbol{s}} = \boldsymbol{x}^{\boldsymbol{s}} \\
\boldsymbol{\alpha}^{\boldsymbol{s}} + \mu^{\boldsymbol{s}} \mathbf{e} \le \mathbf{0}
\end{array}\right\} & \forall \boldsymbol{s} \in \mathcal{S}',
\end{array}
\qquad (\mathcal{CCG}^{\text{master}}_{\text{risk},\Gamma}(\widetilde{\mathcal{S}}))
$$

where $\mathcal{S}' \subseteq \widetilde{\mathcal{S}}^K$. To identify indices $\boldsymbol{s}$ that, given a solution $(\tau, \boldsymbol{\iota})$ to the relaxed master problem, are violated, we solve a *single* feasibility MBLP defined through

$$
\begin{array}{lll}
\min & \theta & \\
\text{s.t.} & \theta \in \mathbb{R},\ \boldsymbol{\epsilon} \in \mathbb{R}^K_{+},\ \boldsymbol{u}^{\boldsymbol{x}} \in \mathcal{U}^0 \ \ \forall \boldsymbol{x} \in \mathcal{R},\ \boldsymbol{s} \in \widetilde{\mathcal{S}}^K & \\
& \theta \ge (\boldsymbol{u}^{\boldsymbol{x}})^\top \boldsymbol{x} \quad \forall \boldsymbol{x} \in \mathcal{R} & \\
& \left.\begin{array}{l}
(\boldsymbol{u}^{\boldsymbol{x}})^\top (\boldsymbol{x}^{\boldsymbol{\iota}_k} - \boldsymbol{x}^{\boldsymbol{\iota}'_k}) + \boldsymbol{\epsilon}_\kappa \ge -M(1 - \boldsymbol{s}_\kappa) \\
(\boldsymbol{u}^{\boldsymbol{x}})^\top (\boldsymbol{x}^{\boldsymbol{\iota}_k} - \boldsymbol{x}^{\boldsymbol{\iota}'_k}) - \boldsymbol{\epsilon}_\kappa \le M(\boldsymbol{s}_\kappa + 1) \\
\displaystyle\sum_{\kappa \in \mathcal{K}} \boldsymbol{\epsilon}_\kappa \le \Gamma
\end{array}\right\} & \forall \kappa \in \mathcal{K} \quad \forall \boldsymbol{x} \in \mathcal{R}.
\end{array}
\qquad (\mathcal{CCG}^{\text{feas}}_{\text{risk},\Gamma}(\boldsymbol{\iota}))
$$

Similar to the noiseless setting, the master and feasibility subproblems can be used to solve Problem $(\mathcal{P}^{\Gamma,K}_{\text{off,risk}})$ iteratively, as summarized in the following Theorem.

Theorem 9. *Algorithm 2 with $(\mathcal{P}^{\Gamma,K}_{\text{off,risk}})$ in place of $(\mathcal{P}^{K}_{\text{off,risk}})$, $(\mathcal{CCG}^{\text{master}}_{\text{risk},\Gamma}(\widetilde{\mathcal{S}}))$ in place of $(\mathcal{CCG}^{\text{master}}_{\text{risk}}(\mathcal{S}'))$, and $(\mathcal{CCG}^{\text{feas}}_{\text{risk},\Gamma}(\boldsymbol{\iota}))$ in place of $(\mathcal{CCG}^{\text{feas}}_{\text{risk}}(\boldsymbol{\iota}))$ terminates in a final number of steps with a feasible solution to Problem $(\mathcal{P}^{\Gamma,K}_{\text{off,risk}})$. The objective value attained by this solution is within δ of the optimal objective value of the problem.*

6.4. Offline Regret Averse Active Preference Elicitation under Inconsistencies

The offline regret averse active preference elicitation problem with inconsistencies is expressible as the following two-stage robust optimization problem with decision-dependent information discovery

$$
\underset{\boldsymbol{\iota} \in \mathcal{C}^K}{\text{minimize}} \quad \max_{\boldsymbol{s} \in \mathcal{S}_\Gamma(\boldsymbol{\iota})} \quad \min_{\boldsymbol{x} \in \mathcal{R}} \quad \max_{\boldsymbol{u} \in \mathcal{U}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s})} \quad \left\{ \max_{\boldsymbol{x}' \in \mathcal{R}} \ \boldsymbol{u}^\top \boldsymbol{x}' - \boldsymbol{u}^\top \boldsymbol{x} \right\}. \qquad (\mathcal{P}^{\Gamma,K}_{\text{off,regret}})
$$

As before, we can simplify the problem above by eliminating the dependence of $\mathcal{S}_\Gamma(\boldsymbol{\iota})$ on $\boldsymbol{\iota}$ and by noting that "indifferent" answers can always be omitted.

In the case of a moderate number of queries K, we can leverage the following proposition to solve Problem ($\mathcal{P}^{\Gamma,K}_{\text{off,regret}}$) by enumeration using the regret averse recommendation problem with inconsistencies ($\mathcal{R}^{\Gamma}_{\text{regret}}$), in a way that parallels Algorithm 3.

Proposition 6. *For any fixed $\boldsymbol{\iota} \in \mathcal{C}^K$ and $\boldsymbol{s} \in \widetilde{\mathcal{S}}^K$, the regret averse recommendation problem with inconsistencies ($\mathcal{R}^{\Gamma}_{\text{regret}}$) is equivalent to the minimization problem*

$$
\begin{array}{ll}
\text{minimize} & \theta \\
\text{subject to} & \theta \in \mathbb{R},\ \boldsymbol{x} \in \mathcal{R} \\
& \left.\begin{array}{l}
\boldsymbol{\alpha}^{\boldsymbol{x}'} \in \mathbb{R}^K_-,\ \boldsymbol{\beta}^{\boldsymbol{x}'} \in \mathbb{R}^M_-,\ \mu^{\boldsymbol{x}'} \in \mathbb{R}_+ \\
\theta \geq \boldsymbol{b}^\top \boldsymbol{\beta}^{\boldsymbol{x}'} + \Gamma \mu^{\boldsymbol{x}'} \\
\displaystyle\sum_{\kappa \in \mathcal{K}} \boldsymbol{s}_\kappa \left(\boldsymbol{x}^{\boldsymbol{\iota}^\kappa_1} - \boldsymbol{x}^{\boldsymbol{\iota}^\kappa_2}\right) \boldsymbol{\alpha}^{\boldsymbol{x}'}_\kappa + \boldsymbol{B}^\top \boldsymbol{\beta}^{\boldsymbol{x}'} = \boldsymbol{x}' - \boldsymbol{x} \\
\boldsymbol{\alpha}^{\boldsymbol{x}'} + \mu^{\boldsymbol{x}'} \mathbf{e} \geq \mathbf{0}
\end{array}\right\} \quad \forall \boldsymbol{x}' \in \mathcal{R},
\end{array}
\tag{13}
$$

whose size is polynomial in the size of the input. Problem (13) *is an MBLP if $\mathcal{R}$ has fixed finite cardinality.*

When K is larger than 2 or 3, solving Problem ($\mathcal{P}^{\Gamma,K}_{\text{off,regret}}$) by enumeration is computationally prohibitive. In such settings, the following lemma and theorem can be leveraged to solve ($\mathcal{P}^{\Gamma,K}_{\text{off,regret}}$) as an MBLP.

Lemma 8. *Problem ($\mathcal{P}^{\Gamma,K}_{\text{off,regret}}$) is equivalent to the following finite program*

$$
\begin{array}{ll}
\text{minimize} & \tau \\
\text{subject to} & \tau \in \mathbb{R},\ \boldsymbol{\iota} \in \mathcal{C}^K,\ \boldsymbol{x}^{\boldsymbol{s}} \in \mathcal{R},\ \boldsymbol{s} \in \widetilde{\mathcal{S}}^K \\
& \left.\begin{array}{l}
\boldsymbol{\alpha}^{(\boldsymbol{x}',\boldsymbol{s})} \in \mathbb{R}^K_-,\ \boldsymbol{\beta}^{(\boldsymbol{x}',\boldsymbol{s})} \in \mathbb{R}^M_-,\ \mu^{(\boldsymbol{x}',\boldsymbol{s})} \in \mathbb{R}_+ \\
\tau \geq \boldsymbol{b}^\top \boldsymbol{\beta}^{(\boldsymbol{x}',\boldsymbol{s})} + \Gamma \mu^{(\boldsymbol{x}',\boldsymbol{s})} \\
\displaystyle\sum_{\kappa \in \mathcal{K}} \boldsymbol{s}_\kappa \left(\boldsymbol{x}^{\boldsymbol{\iota}^\kappa_1} - \boldsymbol{x}^{\boldsymbol{\iota}^\kappa_2}\right) \boldsymbol{\alpha}^{(\boldsymbol{x}',\boldsymbol{s})}_\kappa + \boldsymbol{B}^\top \boldsymbol{\beta}^{(\boldsymbol{x}',\boldsymbol{s})} = \boldsymbol{x}' - \boldsymbol{x}^{\boldsymbol{s}} \\
\boldsymbol{\alpha}^{(\boldsymbol{x}',\boldsymbol{s})} + \mu^{(\boldsymbol{x}',\boldsymbol{s})} \mathbf{e} \geq \mathbf{0}
\end{array}\right\} \quad \forall \boldsymbol{s} \in \widetilde{\mathcal{S}}^K,\ \forall \boldsymbol{x}' \in \mathcal{R},
\end{array}
\tag{14}
$$

where $\boldsymbol{\iota}$ denotes the queries to make and $\boldsymbol{x}^{\boldsymbol{s}}$ the items to recommend in response scenario $\boldsymbol{s} \in \widetilde{\mathcal{S}}^K$.

Theorem 10. *Problem* (14) *is equivalent to the following mixed-binary linear program*

$$
\begin{array}{lll}
\text{minimize} & \tau & \\
\text{subject to} & \tau \in \mathbb{R},\, \boldsymbol{v}^\kappa,\, \boldsymbol{w}^\kappa \in \{0,1\}^I,\, \kappa \in \mathcal{K} & \\
& \boldsymbol{x}^{\boldsymbol{s}} \in \mathcal{R} \quad \forall \boldsymbol{s} \in \widetilde{\mathcal{S}}^K & \\
& \boldsymbol{\alpha}^{(\boldsymbol{x}',\boldsymbol{s})} \in \mathbb{R}^K_-,\, \boldsymbol{\beta}^{(\boldsymbol{x}',\boldsymbol{s})} \in \mathbb{R}^M_-,\, \mu^{(\boldsymbol{x}',\boldsymbol{s})} \in \mathbb{R}_+ \quad \forall \boldsymbol{s} \in \widetilde{\mathcal{S}}^K,\, \boldsymbol{x}' \in \mathcal{R} & \\
& \overline{\boldsymbol{v}}^{(\boldsymbol{x}',\boldsymbol{s})\kappa},\, \overline{\boldsymbol{w}}^{(\boldsymbol{x}',\boldsymbol{s})\kappa} \in \mathbb{R}^I_- \quad \forall \boldsymbol{s} \in \widetilde{\mathcal{S}}^K,\, \boldsymbol{x}' \in \mathcal{R},\, \kappa \in \mathcal{K} & \\
& \left.\begin{array}{l}
\tau \ \geq \ \boldsymbol{b}^\top \boldsymbol{\beta}^{(\boldsymbol{x}',\boldsymbol{s})} + \Gamma \mu^{(\boldsymbol{x}',\boldsymbol{s})} \\
\displaystyle\sum_{i \in \mathcal{I}} \boldsymbol{x}^i \sum_{\kappa \in \mathcal{K}} \boldsymbol{s}_\kappa \left(\overline{\boldsymbol{v}}_i^{(\boldsymbol{x}',\boldsymbol{s})\kappa} - \overline{\boldsymbol{w}}_i^{(\boldsymbol{x}',\boldsymbol{s})\kappa} \right) + \boldsymbol{B}^\top \boldsymbol{\beta}^{(\boldsymbol{x}',\boldsymbol{s})} \ = \ \boldsymbol{x}' - \boldsymbol{x}^{\boldsymbol{s}} \\
\boldsymbol{\alpha}^{(\boldsymbol{x}',\boldsymbol{s})} + \mu^{(\boldsymbol{x}',\boldsymbol{s})} \mathbf{e} \ \geq \ \mathbf{0}
\end{array}\right\} & \begin{array}{c} \forall \boldsymbol{s} \in \widetilde{\mathcal{S}}^K, \\ \boldsymbol{x}' \in \mathcal{R} \end{array} \\
& \mathbf{e}^\top \boldsymbol{v}^\kappa = 1,\, \mathbf{e}^\top \boldsymbol{w}^\kappa = 1,\, 1 - \boldsymbol{w}_i^\kappa \ \geq \displaystyle\sum_{i' : i' \geq i} \boldsymbol{v}_{i'}^\kappa \quad \forall i \in \mathcal{I},\, \forall \kappa \in \mathcal{K} & \\
& \left.\begin{array}{l}
\overline{\boldsymbol{v}}^{(\boldsymbol{x}',\boldsymbol{s})\kappa} \ \geq \ -M \boldsymbol{v}^\kappa,\, \overline{\boldsymbol{v}}^{(\boldsymbol{x}',\boldsymbol{s})\kappa} \ \geq \ \boldsymbol{\alpha}_\kappa^{(\boldsymbol{x}',\boldsymbol{s})} \mathbf{e} \\
\overline{\boldsymbol{v}}^{(\boldsymbol{x}',\boldsymbol{s})\kappa} \ \leq \ \boldsymbol{\alpha}_\kappa^{(\boldsymbol{x}',\boldsymbol{s})} \mathbf{e} + M(\mathbf{e} - \boldsymbol{v}^\kappa) \\
\overline{\boldsymbol{w}}^{(\boldsymbol{x}',\boldsymbol{s})\kappa} \ \geq \ -M \boldsymbol{w}^\kappa,\, \overline{\boldsymbol{w}}^{(\boldsymbol{x}',\boldsymbol{s})\kappa} \ \geq \ \boldsymbol{\alpha}_\kappa^{(\boldsymbol{x}',\boldsymbol{s})} \mathbf{e} \\
\overline{\boldsymbol{w}}^{(\boldsymbol{x}',\boldsymbol{s})\kappa} \ \leq \ \boldsymbol{\alpha}_\kappa^{(\boldsymbol{x}',\boldsymbol{s})} \mathbf{e} + M(\mathbf{e} - \boldsymbol{w}^\kappa)
\end{array}\right\} & \begin{array}{c} \forall \boldsymbol{s} \in \widetilde{\mathcal{S}}^K,\, \boldsymbol{x}' \in \mathcal{R}, \\ \kappa \in \mathcal{K}, \end{array}
\end{array}
\tag{15}
$$

where M is a "big-M" constant. In particular, given an optimal solution $(\tau, \boldsymbol{v}, \boldsymbol{w}, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{x}, \{\mu^{\boldsymbol{s}}\}_{\boldsymbol{s} \in \widetilde{\mathcal{S}}^K}, \overline{\boldsymbol{v}}, \overline{\boldsymbol{w}})$ to Problem (15)*, an optimal set of queries for the risk averse offline preference elicitation problem is given by*

$$
\boldsymbol{\iota}_1^\kappa = \sum_{i \in \mathcal{I}} i \cdot \mathbb{I}(\boldsymbol{v}_i^\kappa = 1) \qquad \textit{and} \qquad \boldsymbol{\iota}_2^\kappa = \sum_{i \in \mathcal{I}} i \cdot \mathbb{I}(\boldsymbol{w}_i^\kappa = 1), \quad \kappa \in \mathcal{K}.
$$

Problem (15) can be solved using a column-and-constraint generation procedure similar to that proposed in Section 5.4 for the case without inconsistencies. Here, we only provide the master problem, the feasibility subproblem, and the algorithm. The relaxed master problem is

$$
\begin{array}{lll}
\text{minimize} & \tau & \\
\text{subject to} & \tau \in \mathbb{R},\, \boldsymbol{\iota} \in \mathcal{C}^K,\, \boldsymbol{x}^{\boldsymbol{s}} \in \mathcal{R},\, \boldsymbol{s} \in \widetilde{\mathcal{S}}^K & \\
& \left.\begin{array}{l}
\boldsymbol{\alpha}^{(\boldsymbol{x}',\boldsymbol{s})} \in \mathbb{R}^K_-,\, \boldsymbol{\beta}^{(\boldsymbol{x}',\boldsymbol{s})} \in \mathbb{R}^M_-,\, \mu^{(\boldsymbol{x}',\boldsymbol{s})} \in \mathbb{R}_+ \\
\tau \ \geq \ \boldsymbol{b}^\top \boldsymbol{\beta}^{(\boldsymbol{x}',\boldsymbol{s})} + \Gamma \mu^{(\boldsymbol{x}',\boldsymbol{s})} \in \mathbb{R}_+ \\
\displaystyle\sum_{\kappa \in \mathcal{K}} \boldsymbol{s}_\kappa \left(\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa} \right) \boldsymbol{\alpha}_\kappa^{(\boldsymbol{x}',\boldsymbol{s})} + \boldsymbol{B}^\top \boldsymbol{\beta}^{(\boldsymbol{x}',\boldsymbol{s})} \ = \ \boldsymbol{x}^{\boldsymbol{s}} - \boldsymbol{x}' \\
\boldsymbol{\alpha}^{(\boldsymbol{x}',\boldsymbol{s})} + \mu^{(\boldsymbol{x}',\boldsymbol{s})} \mathbf{e} \ \geq \ \mathbf{0}
\end{array}\right\} & \begin{array}{c} \forall \boldsymbol{s} \in \mathcal{S}', \\ \boldsymbol{x}' \in \mathcal{R}', \end{array}
\end{array}
\tag{$\mathcal{CCG}^{\text{master}}_{\text{regret},\Gamma}(\widetilde{\mathcal{S}})$}
$$

where $\mathcal{S}' \subseteq \widetilde{\mathcal{S}}^K$. To identify indices $\boldsymbol{s}$ that, given a solution $(\tau, \boldsymbol{\iota})$ to the relaxed master problem, are violated, we solve a *single* feasibility MBLP defined through

$$
\begin{array}{ll}
\min & \theta \\
\text{s.t.} & \theta \in \mathbb{R},\ \boldsymbol{\epsilon} \in \mathbb{R}^K_+,\ \boldsymbol{u}^{\boldsymbol{x}} \in \mathcal{U}^0 \quad \forall \boldsymbol{x} \in \mathcal{R},\ \boldsymbol{s} \in \overline{\mathcal{S}}^K \\
& \theta \geq (\boldsymbol{u}^{\boldsymbol{x}})^\top \boldsymbol{x} \quad \forall \boldsymbol{x} \in \mathcal{R} \\
& \left.\begin{array}{l}
(\boldsymbol{u}^{\boldsymbol{x}})^\top (\boldsymbol{x}^{\boldsymbol{\iota}_k} - \boldsymbol{x}^{\boldsymbol{\iota}'_k}) + \boldsymbol{\epsilon}_\kappa \geq -M(1 - \boldsymbol{s}_\kappa) \\
(\boldsymbol{u}^{\boldsymbol{x}})^\top (\boldsymbol{x}^{\boldsymbol{\iota}_k} - \boldsymbol{x}^{\boldsymbol{\iota}'_k}) - \boldsymbol{\epsilon}_\kappa \leq M(\boldsymbol{s}_\kappa + 1) \\
\sum_{\kappa \in \mathcal{K}} \boldsymbol{\epsilon}_\kappa \leq \Gamma
\end{array}\right\} \quad \forall \kappa \in \mathcal{K} \quad \forall \boldsymbol{x} \in \mathcal{R}.
\end{array}
\qquad (\mathcal{CCG}^{\text{feas}}_{\text{regret},\Gamma}(\boldsymbol{\iota}))
$$

Theorem 11. *Algorithm 4 with* $(\mathcal{P}^{\Gamma,K}_{\text{off,regret}})$ *in place of* $(\mathcal{P}^{K}_{\text{off,regret}})$, $(\mathcal{CCG}^{\text{master}}_{\text{regret},\Gamma}(\widetilde{\mathcal{S}}))$ *in place of* $(\mathcal{CCG}^{\text{master}}_{\text{regret},\Gamma}(\widetilde{\mathcal{S}}))$, *and* $(\mathcal{CCG}^{\text{feas}}_{\text{regret},\Gamma}(\boldsymbol{\iota}))$ *in place of* $(\mathcal{CCG}^{\text{feas}}_{\text{regret}}(\boldsymbol{\iota}))$ *terminates in a final number of steps with a feasible solution to Problem* $(\mathcal{P}^{\Gamma,K}_{\text{off,regret}})$. *The objective value attained by this solution is within* δ *of the optimal objective value of the problem.*

7. Speed-up Strategies

In this section, we propose two complementary strategies for speeding up solution of the mixed-binary linear programming reformulations of the active preference elicitation problems (Problems (6),(9),(12), and (15)). The first one aims to break the symmetry in the problem and the second adds warm starts.

7.1. Symmetry Breaking

In the finite programming reformulations (5),(8),(11), and (14) of Problems $(\mathcal{P}^{K}_{\text{off,risk}})$,$(\mathcal{P}^{K}_{\text{off,regret}})$,$(\mathcal{P}^{\Gamma,K}_{\text{off,risk}})$, and $(\mathcal{P}^{\Gamma,K}_{\text{off,regret}})$, every permutation of the queries $\boldsymbol{\iota}$ will yield another solution with the same objective function value. Indeed, since all queries are asked simultaneously, the index of a query does not impact the answer of the user to the query. Correspondingly, Problems (6),(9),(12), and (15) also present several symmetric solutions: we can permute the indices κ in the pairs $\{(\boldsymbol{v}^\kappa, \boldsymbol{w}^\kappa)\}_{\kappa \in \mathcal{K}}$ to build solutions with the same objective. This symmetry increases the size of the search space and therefore, time is spent visiting solutions which are symmetric to the ones already visited in the search tree. To speed-up solution time, we propose to augment the MBLP formulations with symmetry breaking constraints which eliminate symmetric solutions from the search space while preserving at least one solution from each equivalence class.

To build the symmetry breaking constraints, note that each comparison is uniquely defined by the binary vector

$$\boldsymbol{y}^\kappa = \boldsymbol{v}^\kappa + \boldsymbol{w}^\kappa, \tag{16}$$

which has exactly two nonzero elements – corresponding to the items used in the comparison. To break this symmetry, we require that the binary vectors $\{\boldsymbol{y}^\kappa\}_{\kappa\in\mathcal{K}}$ be *lexicographically ordered*: if $\boldsymbol{y}_i^\kappa = \boldsymbol{y}_i^{\kappa+1}$ for all $i<j$, and $\boldsymbol{y}_j^\kappa \neq \boldsymbol{y}_j^{\kappa+1}$, then $\boldsymbol{y}_j^\kappa = 0$ and $\boldsymbol{y}_j^{\kappa+1} = 1$. To enforce lexicographic ordering, we introduce, in addition to $\boldsymbol{y}^\kappa$, the binary variables $\boldsymbol{z}^{\kappa\kappa'} \in \{0,1\}^I$, which satisfy $\boldsymbol{z}_i^{\kappa\kappa'} = 1$ if and only if $\boldsymbol{y}_i^\kappa \neq \boldsymbol{y}_i^{\kappa'}$, $i\in\mathcal{I}$. These variables can be uniquely defined using the following linear constraints

$$\left.\begin{aligned}
\boldsymbol{z}^{\kappa\kappa'} &\leq \boldsymbol{y}^\kappa + \boldsymbol{y}^{\kappa'} \\
\boldsymbol{z}^{\kappa\kappa'} &\leq 2 - \boldsymbol{y}^\kappa - \boldsymbol{y}^{\kappa'} \\
\boldsymbol{z}^{\kappa\kappa'} &\geq \boldsymbol{y}^\kappa - \boldsymbol{y}^{\kappa'} \\
\boldsymbol{z}^{\kappa\kappa'} &\geq \boldsymbol{y}^{\kappa'} - \boldsymbol{y}^\kappa
\end{aligned}\right\} \quad \forall \kappa,\ \kappa' \in \mathcal{K} \,:\, \kappa < \kappa'. \tag{17}$$

Lexicographic ordering can then be imposed by adding to all MBLP problems the constraints

$$\boldsymbol{y}_i^{\kappa'} \geq \boldsymbol{y}_i^\kappa - \sum_{i'\in\mathcal{I}:i'<i} \boldsymbol{z}_{i'}^{\kappa\kappa'} \quad \forall i\in\mathcal{I},\ \kappa,\kappa'\in\mathcal{K}:\kappa<\kappa'. \tag{18}$$

Next, note that in the preference elicitation problems, there is no benefit in asking the same query twice: the worst-case utility can only improve by making different queries. Correspondingly, we can add to all MBLP problems the following constraints

$$\sum_{i\in\mathcal{I}} \boldsymbol{z}_i^{\kappa,\kappa'} \geq 1 \quad \forall \kappa,\kappa'\in\mathcal{K} \,:\, \kappa<\kappa', \tag{19}$$

which eliminate solutions that involve asking the same queries twice.

As we will see in Section 8, the symmetry breaking constraints (16)-(18), and (19) translate to a significant reduction in solver times.

7.2. Warm-Starts

Simultaneously to the first speed-up strategy discussed in Section 7.1, solution times can be further reduced by employing warm starts. In particular, solutions obtained for lower values of K (smaller numbers of queries) can be used to warm start problems with more queries. In this section, we describe this procedure. We detail how to build a feasible warm start for Problem (6) with $K+1$ queries from a feasible solution to Problem (6) with K queries. Recall that a warm start need only fix solutions for the binary decision variables in the problem. The approach naturally also applies to MBLPs (9),(12), and (15).

Let $(\{\tau, \boldsymbol{v}, \boldsymbol{w}, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{x}, \overline{\boldsymbol{v}}, \overline{\boldsymbol{w}})$ be feasible in Problem (6) with K queries. To generate a feasible solution to Problem (6) with $K+1$ queries, we proceed as follows. First, we generate a new query $\boldsymbol{\iota}^{K+1}$ randomly from the set

$$\mathcal{C} \setminus \left\{ \boldsymbol{\iota} \in \mathcal{C} \,:\, \exists \kappa \in \mathcal{K} \,:\, \boldsymbol{\iota}_1^\kappa = \sum_{i \in \mathcal{I}} i \cdot \mathbb{I}\left(\boldsymbol{v}_i^\kappa = 1\right) \text{ and } \boldsymbol{\iota}_2^\kappa = \sum_{i \in \mathcal{I}} i \cdot \mathbb{I}\left(\boldsymbol{w}_i^\kappa = 1\right) \right\}$$

and define $\tilde{\boldsymbol{v}}$ and $\tilde{\boldsymbol{w}}$ through $\tilde{\boldsymbol{v}}^\kappa := \boldsymbol{v}^\kappa$, $\tilde{\boldsymbol{w}}^\kappa := \boldsymbol{w}^\kappa$, $\kappa \in \mathcal{K}$, $\tilde{\boldsymbol{v}}_i^{K+1} := \mathbb{I}\left(\boldsymbol{\iota}_1^{K+1} = i\right)$, and $\tilde{\boldsymbol{w}}_i^{K+1} := \mathbb{I}\left(\boldsymbol{\iota}_2^{K+1} = i\right)$ for each $i \in \mathcal{I}$. Subsequently, for each $\boldsymbol{s} \in \widetilde{\mathcal{S}}^{K+1}$, we define $\tilde{\boldsymbol{x}}^{\boldsymbol{s}} := \boldsymbol{x}^{(\boldsymbol{s}_1, \ldots, \boldsymbol{s}_K)}$. Then, $(\tilde{\boldsymbol{v}}, \tilde{\boldsymbol{w}}, \tilde{\boldsymbol{x}})$ constitutes a warm start for Problem (6) with $K+1$ queries. The solution hereby constructed asks one additional (random) query and subsequently ignores the answer to the query in the choice of recommendation. It can be shown that it attains an objective value in Problem (6) with $K+1$ queries at least as high as that attained by $(\tau, \boldsymbol{v}, \boldsymbol{w}, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{x}, \overline{\boldsymbol{v}}, \overline{\boldsymbol{w}})$ in Problem (6) with K queries. Note that if warm starts are combined with symmetry breaking constraints, then the solution constructed needs to be permuted so as to satisfy lexicographical ordering to ensure it satisfies the lexicographic constraints. This procedure is detailed in Algorithm 5. As we will see in our numerical results, see Section 8, using warm starts in this fashion results in significant speed-ups.

8. Numerical Results

In this section, we investigate the performance of our approach on both synthetic and real-world instances of the offline and online max-min utility and min-max regret active preference elicitation problems introduced in this paper. This section is organized as follows. In Sections 8.1 and 8.2, we describe the datasets that we employ in our experiments. The experimental setup is described in Section 8.3. In Sections 8.4 and 8.5, we present our numerical results on the synthetic and real datasets, respectively.

8.1. Description of the Synthetic Datasets

We generate synthetic datasets with varying numbers of items $I \in \{7, 10, 20, 30, 40, 60\}$ and features $J \in \{3, 5, 10, 20\}$. Items $\boldsymbol{x}^i$ are drawn uniformly at random from the J-dimensional sphere of radius ten. Throughout our synthetic data experiments, we assume $\mathcal{U}^0 = [-1, 1]^J$.

8.2. Description of the Real Dataset

We use real data from the HMIS database obtained from Ian De Jong[8] as part of a working group called "Youth Homelessness Data, Policy, Research." The dataset tracks 10,922 homeless youth from 16 communities across the United States and presents 3,474 housing assignments. Each row in the dataset corresponds to a youth that was waitlisted for housing. There are two types of housing resources: rapid rehousing (RRH) and permanent supportive housing (PSH). Some individuals received services only (SO). For each individual, the dataset presents their race/ethnicity, their TAY VI-SPDAT[9] score, the housing resource they were allocated, and the outcome of the intervention (whether they successfully exited homelessness for one year or longer). We refer the reader to Chan et al. (2017) for more information on this dataset. We classify individuals into 31 classes based on their TAY VI-SPDAT score (1 through 16) and race ("White or Other" and "Black or Hispanic"). We use logistic regression to learn the success probabilities for each class of individuals and for each of the three interventions (PSH, RRH, SO).

Using this data, we build a dataset of $I = 50$ policies and their $J = 21$ outcomes/characteristics through simulation. We now detail this procedure which proceeds in two steps: the generation of a policy and the simulation of the performance of that policy based on historical data.

To generate policies, we propose to view the system as a multi-class multi-server queuing system under a first come, first served (FCFS) priority rule. In particular, we think of the system as a bipartite graph with one side of the graph collecting the classes (queues) of individuals of different types and the other side collecting the servers where housing resources of different types are procured, see Figure 3. Thus, each queue collects all individuals from the same class and each server procures housing resources of a single type. A link between a server and a class in the bipartite graph indicates that individuals of this class can be matched to resources procured by the server. When a resource is procured at a server, it is offered to the first person to have arrived in any one of the queues eligible for service from this server (i.e., connected to this server). If an individual is not matched to any resource, they are assumed to receive the SO intervention. A *housing allocation policy* can then be uniquely defined through the topology of this bipartite graph. We generate a set of $I = 50$ topologies as follows. The first topology corresponds to the first come, first served priority discipline (i.e., a complete bipartite graph). The second topology corresponds to the current allocation policy: individuals that score high (TAY VI-SPDAT greater or equal to 8) are waitlisted for PSH; individuals that score medium (TAY VI-SPDAT in the range 4 to 7) are waitlisted for RRH; others receive service only. The

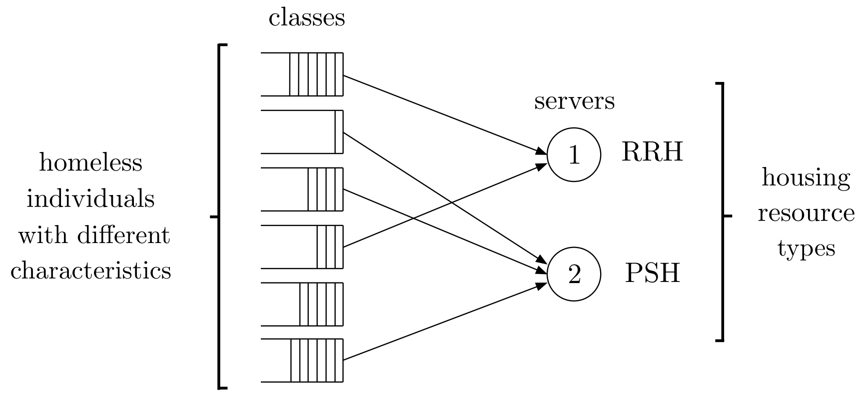

Figure 3 Model of the housing allocation system as a bipartite graph. The policy associated with this bipartite graph is "interpretable" in the sense that each queue is only matched to one server.

remaining topologies are generated at random. Half of the policies are *interpretable* in the sense that each class can only be connected to one server –this ensures that individuals can be guided toward the appropriate intervention upon entering the system. Whether a policy can use race or not as feature to assign housing is a Bernoulli random variable. Topologies that comply with the interpretability and the use or not of protected characteristics are then generated uniformly at random.

Once these 50 topologies have been generated, we simulate the historical performance of the corresponding policy. In particular, for each topology, we simulate the interventions that would have been received by each individual under this alternative policy and record the following attributes: *a)* the probability that any given individual exits homelessness, overall, by race, by score range (hi, medium, low), and by race and score range combined; *b)* the average treatment effect, by race, by score range (hi, medium), and by race and score range combined; and *c)* whether the policy was interpretable, giving us a total of $J=21$ features. For the experiments on real data, we assume $\mathcal{U}^0 = \{\boldsymbol{u} \in \mathbb{R}^J_+ \,:\, \mathbf{e}^\top \boldsymbol{u} = 1\}$ in the spirit of conjoint analysis, see Example 3.

8.3. Experimental Setup

We solve the offline elicitation problems ($\mathcal{P}^K_{\text{off,risk}}$), ($\mathcal{P}^K_{\text{off,regret}}$), ($\mathcal{P}^{\Gamma,K}_{\text{off,risk}}$), and ($\mathcal{P}^{\Gamma,K}_{\text{off,regret}}$) using the column-and-constraint generation procedures discussed in Sections 3.4, 5.4, 6.3, and 6.4, respectively. The tolerance δ used in the column-and-constraint generation algorithm is 1×10^{-3}. Unless explicitly stated otherwise, we augment all our formulations with the symmetry breaking constraints proposed in Section 7.1. To speed-up computation further, 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.4. This strategy enables us to run a far larger number of experiments in a given time to showcase the performance of our approach in a variety of settings. Without this conservative heuristic, the quality of the solutions obtained by our approach would be higher than those reported in these experiments.

In the offline setting, we compare our elicitation approaches, which we refer to as `E-OFF-MMU` and `E-OFF-MMR` for max-min utility and min-max regret cases, respectively, to random elicitation, which we denote by `E-RAND`. For any fixed K, `E-RAND` selects K queries at random. To the best of our knowledge, no other solution approach exists in the literature for identifying K queries to ask at once. For `E-RAND`, we sample 50 sets of K random queries. We evaluate the worst-case utility (resp. worst-case regret) of any given choice of queries $\boldsymbol{\iota}^\star$, which we denote by $u_{\text{wc}}(\boldsymbol{\iota}^\star)$ (resp. $r_{\text{wc}}(\boldsymbol{\iota}^\star)$), by fixing $\boldsymbol{\iota} = \boldsymbol{\iota}^\star$ in Problem (6) (resp. (9)) for the case without inconsistencies and in Problem (12) (resp. (15)) for the case with inconsistencies. To speed-up this objective function evaluation step, we employ a variant of the column-and-constraint generation procedure presented in this paper. We omit the details in the interest of space. We investigate the role played by all speed-up strategies and the suboptimality of the conservative heuristic approach at the end of Section 8.4.

In the online setting, we compare our conservative solution approaches, which we refer to as `E-ON-MMU` and `E-ON-MMR` for the max-min utility and min-max regret cases, respectively, to various heuristics from the literature. These approaches combine different elicitation and recommendation strategies. In terms of elicitation, we investigate random query selection (denoted by `E-RAND`) and query selection using the analytic center approach (denoted by `E-AC`) as in Bertsimas and O'Hair (2013) (which is reminiscent of the elicitation approaches from Toubia et al. (2003) and Boutilier et al. (2006)). Note that in the case of inconsistencies, we adapt the approach from Bertsimas and O'Hair (2013) to employ our uncertainty set to make for a fair comparison. In terms of recommendation, we investigate analytic center based recommendation (denoted by `R-AC`) (originally proposed by Toubia et al. (2003)) and robust recommendations based on the max-min utility and min-max regret decision criteria (denoted by `R-MMU` and `R-MMR`), which we compute using our reformulations (3), (7), (10), and (13), as appropriate. Thus, in the max-min utility case, we compare our approach, `E-ON-MMU/R-MMU` to: `E-AC/R-AC`, `E-AC/R-MMU`, and `E-RAND/R-MMU`. In the min-max regret case, we compare our approach, `E-ON-MMR/R-MMR` to: `E-AC/R-AC`, `E-AC/R-MMR`, and `E-RAND/R-MMR`. For the online setting, we evaluate the performance of all approaches through simulation using 50 (resp. 13) true utility vectors $\boldsymbol{u}^\star$ for the synthetic (resp. real) dataset. The utility vectors are drawn uniformly at random from the J-dimensional sphere of radius one. For each method (which is a tuple consisting of an elicitation approach and a recommendation approach), we generate queries according to the elicitation approach, generate answers to the queries using $\boldsymbol{u}^\star$, and subsequently make a recommendation according to the recommendation approach. We record the worst-case utility, worst-case regret, and rank of the recommended item.

In line with several recent papers in the robust optimization literature, see e.g., Bandi and Bertsimas (2012), in the case of inconsistent responses, we employ uncertainty sets $\mathcal{U}_\Gamma(\boldsymbol{s}, \boldsymbol{\iota})$, where $\Gamma := \Gamma^{\mathbb{G}}(K)^{1/2} \max_{i\in\mathcal{I}} \|\boldsymbol{x}^i\|_1$. This definition ensures that the interpretation of $\Gamma^{\mathbb{G}}$ and $\boldsymbol{\epsilon}$ remain consistent across experiments and independent of the number of questions K and of the range of $\boldsymbol{u}^\top \boldsymbol{x}^i$, $i \in \mathcal{I}$.

In order to facilitate an interpretation of the performance of our approach, we standardize the utilities (regrets) of the recommended items. In the case of max-min utility, we standardize worst-case utilities such that a standardized worst-case utility equal to zero (resp. one) corresponds to the worst-case utility of the item recommended if no queries (resp. all queries) are asked. Mathematically, the normalized worst-case utility associated with a choice of queries $\boldsymbol{\iota}$ is calculated as

$$u_{\text{nwc}}(\boldsymbol{\iota}) = \frac{u_{\text{wc}}(\boldsymbol{\iota}) - u_{\text{wc}}^0}{u_{\text{wc}}^{\text{all}} - u_{\text{wc}}^0},$$

where u_{wc}^0 (resp. $u_{\text{wc}}^{\text{all}}$) denote the worst-case utility of the item recommended if none (resp. all) queries are asked. Thus, if after asking queries, recommendations are made according to `R-MMU`, then their normalized worst-case utility will be in the range $[0,1]$. On the other hand, if recommendations are made in a way that does not maximize worst-case utility (e.g., according to `R-AC`), then they may have a normalized worst-case utility that is negative. In the case of min-max regret, we standardize the worst-case regrets such that a standardized worst-case regret equal to one (resp. zero) corresponds to the worst-case regret of the item recommended if no queries (resp. all queries) are asked. Mathematically, the normalized worst-case regret is calculated as

$$r_{\text{nwc}}(\boldsymbol{\iota}) = \frac{r_{\text{wc}}(\boldsymbol{\iota}) - r_{\text{wc}}^0}{r_{\text{wc}}^{\text{all}} - r_{\text{wc}}^0},$$

where r_{wc}^0 (resp. $r_{\text{wc}}^{\text{all}}$) denote the worst-case regret of the item recommended if no (resp. all) queries are asked. Thus, if after asking queries, recommendations are made according to `R-MMR`, then their normalized worst-case regret will be in the range $[0,1]$. On the other hand, if recommendations are made in a way that does not minimize worst-case regret, then they may have a normalized worst-case regret that is greater than one.

Across our experiments, we use smaller datasets for the regret averse case than for the risk averse case. The reason is that we benchmark against random elicitation and evaluating the performance of a large number of random queries becomes computationally prohibitive for large datasets in the regret averse case (more so than in the risk averse case).

All of our experiments were performed on a Linux virtual machine (Ubuntu 18.04.1 LTS), with 8GB RAM and two Intel Xeon 2.6GHz cores. All linear optimization problems were solved using Gurobi version 9.0.0. We use Mosek version 9.1.l0 to compute the analytic centers of polyhedra as required by `E-AC` and `R-AC`.

8.4. Numerical Results on Synthetic Data

Offline Preference Elicitation: Optimality and Scalability Results. In our first set of experiments, we evaluate the benefits of our offline approaches relative to random elicitation under both the max-min utility and min-max regret decision criteria. The results are summarized on Figures 4 and 5 for the risk averse case and on Figures 6 and 7 for the regret averse case. Figures 4 and 6 summarize optimality results, while Figures 5 and 7 summarize computational performance. From Figure 4, we see that `E-OFF-MMU` consistently outperforms `E-RAND`. In our experiments, `E-OFF-MMU` outperformed the median query from `E-RAND` by up to 52 percentage points in terms of normalized worst-case utility. Similarly, from Figure 6, we see that `E-OFF-MMR` consistently outperforms `E-RAND`. In our experiments, `E-OFF-MMR` outperformed a median query from `E-RAND` by up to 27 percentage points in terms of worst-case regret. Thus, selecting queries at random almost always results in suboptimal performance. At the same time, from Figure 5, we see that evaluating the performance of a single random query (which is only possible thanks to the techniques we propose in this paper) takes on average 1.22 seconds (resp. 11.68 seconds) for max-min utility (resp. min-max regret) objective. At the same time, given I and K, the total number of possible random queries amounts to $\binom{I(I-1)}{K}$, which is in the order of 7e10 (resp. 1.8e5) for $I = 60$ and $K = 10$ (resp. $I = 20$ and $K = 10$), making it prohibitive to evaluate the performance of all possible random queries to then select the best option. We note that for a given problem size, computing the optimal set of queries in the the regret averse setting is harder than in the risk averse setting.

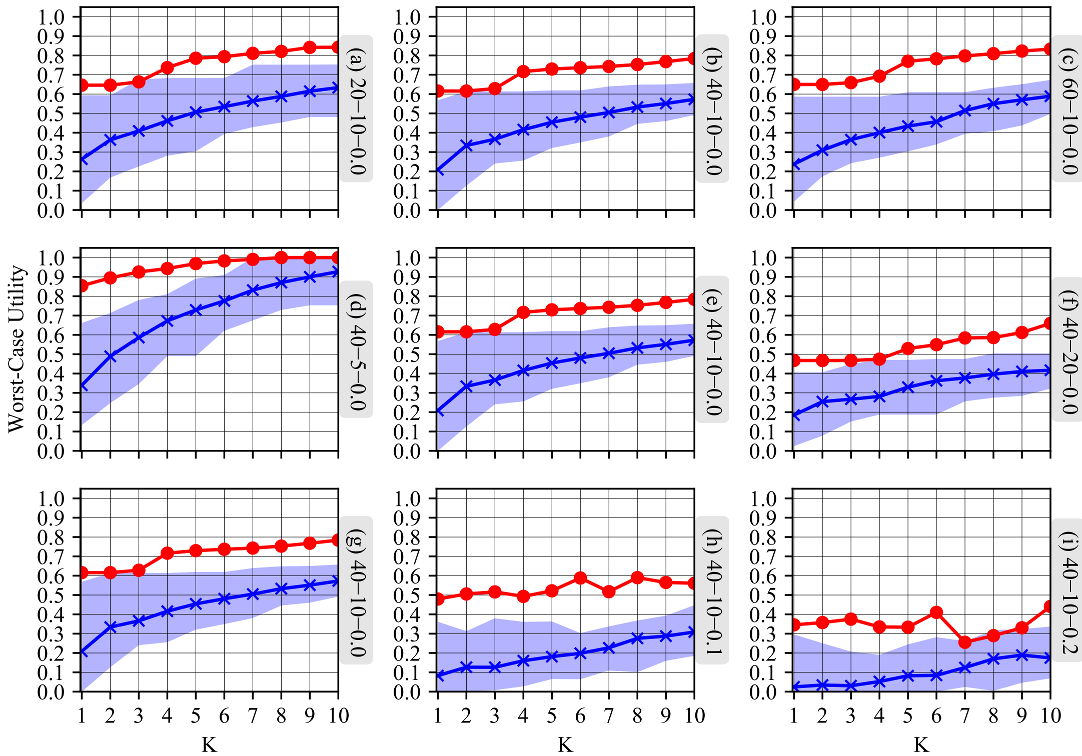

Figure 4 Optimality results for the offline risk averse active preference elicitation problem ($\mathcal{P}^K_{\text{off,risk}}$) on synthetic data. Each label on the right of each facet corresponds to the characteristics of the instance solved ($I - J - \Gamma$). For example, the first facet (subfigure (a)) is labeled 20–10–0.0, indicating an instance with 20 items, 10 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. Approach `E-OFF-MMU` is shown with red dots. The median performance of `E-RAND` across 50 sets of K random queries is shown with blue crosses. The blue shaded region shows the range of performance of `E-RAND` across these 50 sets of queries.

Online Preference Elicitation: Optimality Results. In our second set of experiments, we evaluate the benefits of our online approaches relative to existing approaches from the literature, see Section 8.3 for details. Our results are summarized in Figures 8 and 9 for the max-min utility and min-max regret cases, respectively. From the figures, it is apparent that `E-ON-MMU/R-MMU` and `E-ON-MMR/R-MMR` consistently outperform all other elicitation/recommendation approaches. In particular, in the risk averse case, the greatest improvement of `E-ON-MMU/R-MMU` over `E-AC/R-AC`, `E-AC/R-MMU`, and `E-RAND/R-MMU` is 96, 73, and 76 percentage points, respectively. In the regret averse case, the greatest improvement of `E-ON-MMR/R-MMR` over `E-AC/R-AC`,

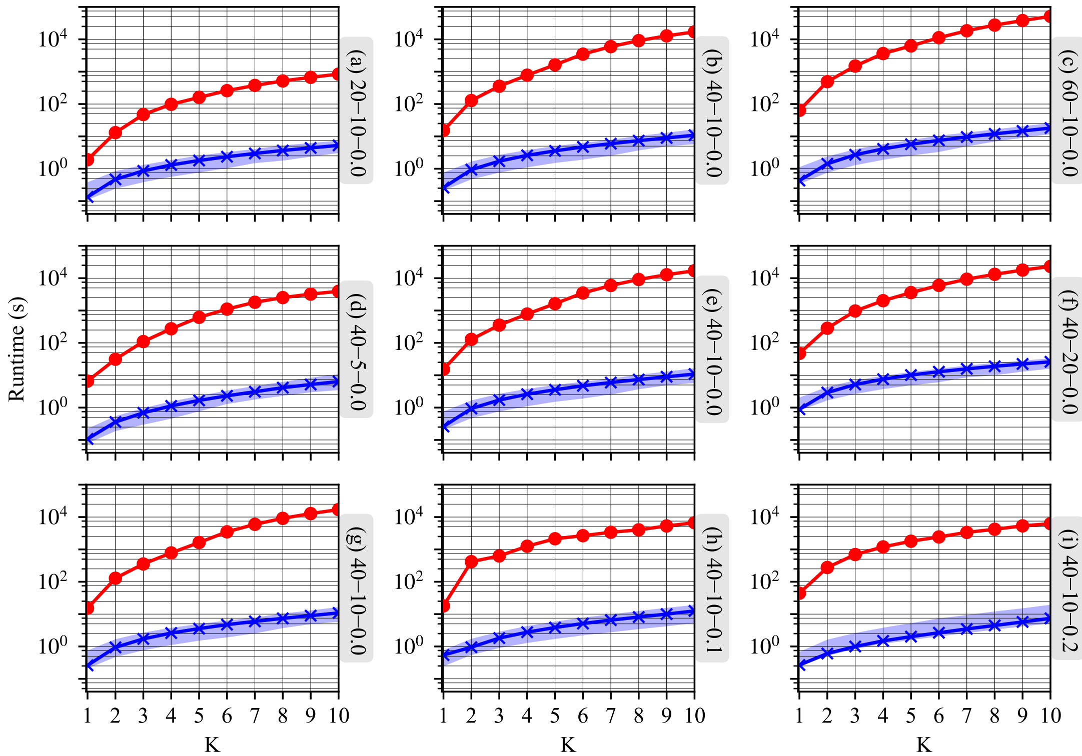

Figure 5 Scalability results for the offline risk averse active preference elicitation problem ($\mathcal{P}^K_{\text{off,risk}}$) on synthetic data. The facet labels, graphs, shapes, lines, and colors have the same interpretation as in Figure 4. In particular, approach `E-OFF-MMU` is shown with red dots and the median (and range of) performance of `E-RAND` across 50 random queries is shown with blue crosses (shade).

`E-AC/R-MMR`, and `E-RAND/R-MMR` is 80, 74, and 73 percentage points, respectively. Moreover, in this online setting, with just a moderate number of queries, a normalized worst-case utility of one was achieved in all cases where $\Gamma = 0$. Similarly, a normalized worst-case regret of zero was reached in most cases without inconsistencies. This indicates that, if chosen strategically, a moderate number of queries is needed to be competitive with the full information case. From the figures, it is also apparent that approach `E-AC/R-AC`, after having asked several queries, may still perform worse that making a recommendation using `R-MMU` or `R-MMR` directly without asking any queries. At the same time, the time needed to compute a query in `E-ON-MMU` and `E-ON-MMR` is 16.48 and 1.57 seconds on average, respectively. Thus, our online elicitation procedure is adequate for use in online (interactive) settings.

Comparison Between Max-Min Utility & Min-Max Regret Solutions. In the first set of experiments, we observed that the max-min utility problem is more scalable than its min-max regret counterpart. In our

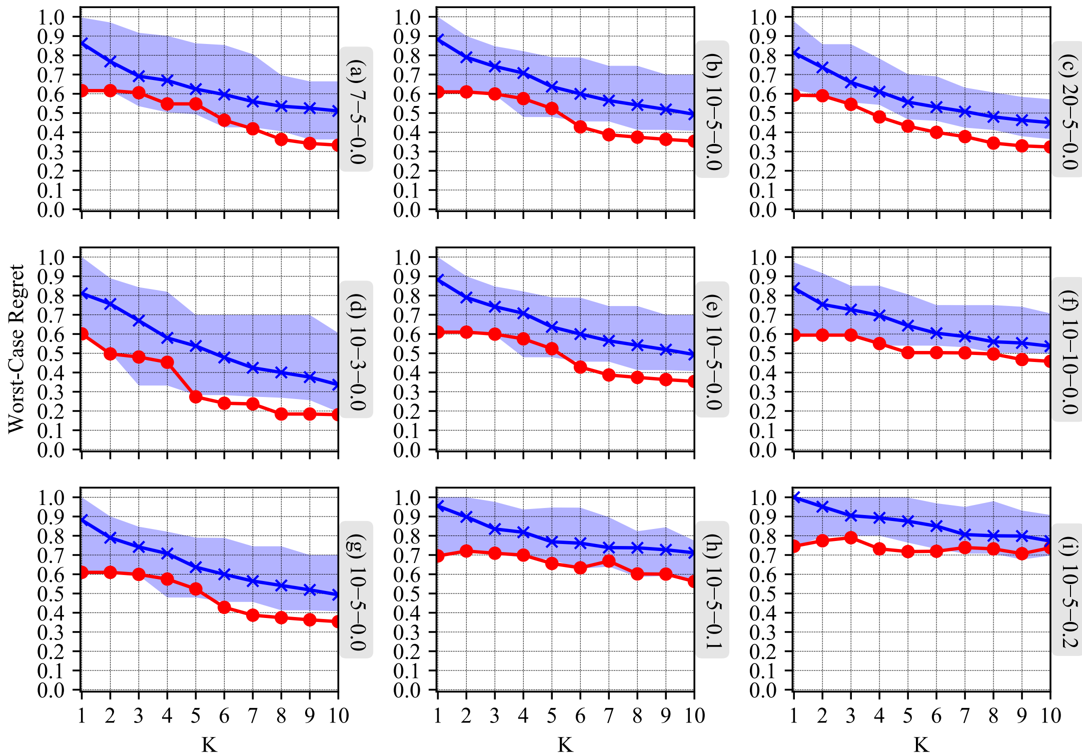

Figure 6 Optimality results for the offline regret averse active preference elicitation problem ($\mathcal{P}^K_{\text{off,regret}}$) on synthetic data. The facet labels, graphs, shapes, lines, and colors have similar interpretation as in Figure 4. In particular, approach `E-OFF-MMR` is shown with red dots and the median (and range of) performance of `E-RAND` across 50 random queries is shown with blue crosses (shade).

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 the worst-case utility and worst-case regret of solutions to Problems ($\mathcal{P}^K_{\text{off,risk}}$) and ($\mathcal{P}^K_{\text{off,regret}}$), 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 employing the min-max regret criterion generally results in solutions that have lower worst-case regret but that do perform more poorly in terms of worst-case utility.

Evaluation of Column-and-Constraint Generation, Symmetry Breaking, & Greedy Heuristic. For our fourth set of experiments, we evaluate the performance of our speed-up strategies. To this end, we solve the offline risk averse active preference elicitation problem ($\mathcal{P}^K_{\text{off,risk}}$) based on three synthetic datasets using four different approaches: the MBLP problem (6), the MBLP problem (6) augmented with symmetry breaking constraints (see Section 7.1), the column-and-constraint generation (CCG) approach from Algorithm 2

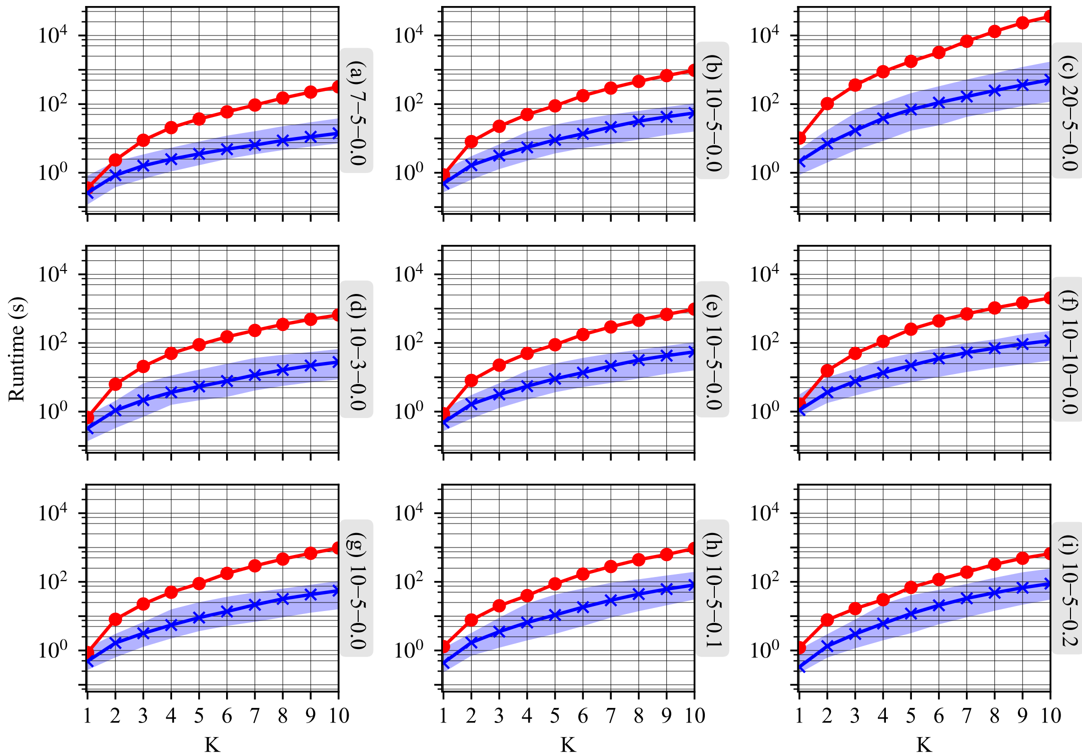

Figure 7 Scalability results for the offline regret averse active preference elicitation problem ($\mathcal{P}^K_{\text{off,regret}}$) on synthetic data. The facet labels, graphs, shapes, lines, and colors have similar interpretation as in Figure 5. In particular, approach `E-OFF-MMR` is shown with red dots and the median (and range of) performance of `E-RAND` across 50 random queries is shown with blue crosses (shade).

augmented with symmetry breaking constraints, and the CCG-based conservative heuristic approach, see Section EC.5.2. The results are provided in Table 2. From the table, it can be seen that the symmetry breaking constraints and the CCG algorithm augmented with symmetry breaking significantly speed-up computation and enable the solution of problems that could not be solved with the MBLP alone. From the table, it also becomes apparent that the CCG-based conservative heuristic returns near-optimal solutions in a fraction of the time, thus scaling to far larger instances.

8.5. Numerical Results on Real Data

We evaluate the benefits of our approach on the real dataset described in Section 8.2. For this dataset, the worst-case utility after asking no queries and after asking all queries are equal and therefore no useful

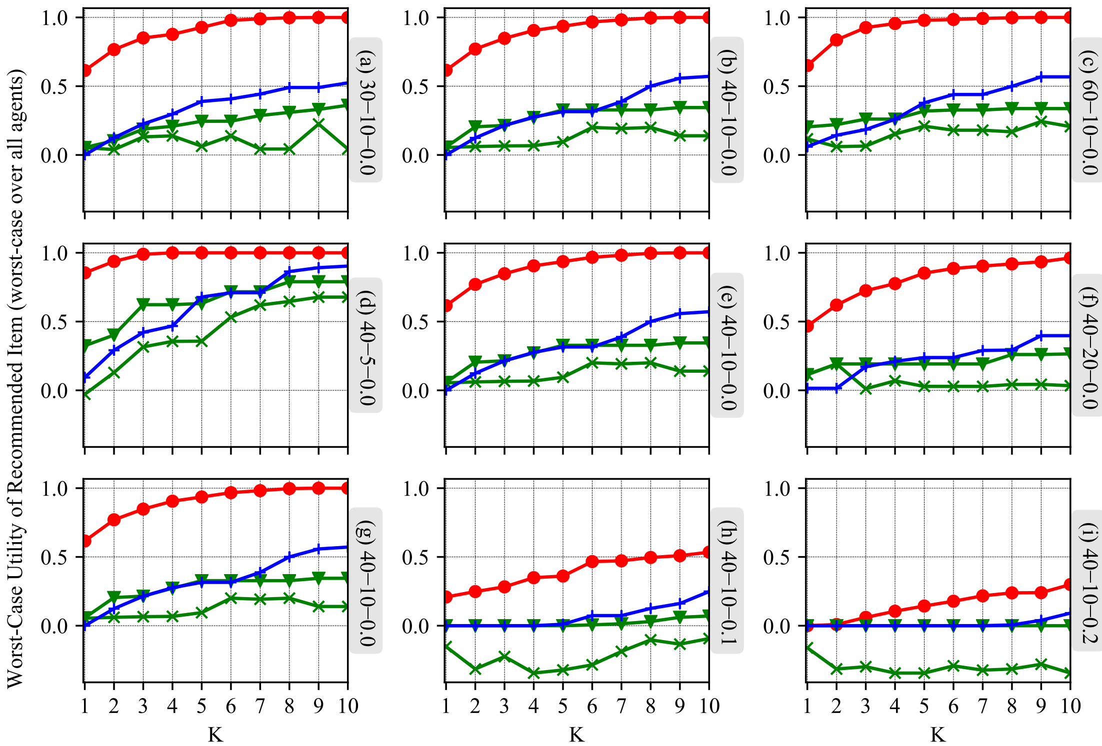

Figure 8 Optimality results for the online risk averse active preference elicitation problem ($\mathcal{P}^K_{\text{on,risk}}$) on synthetic data. The facet labels have the same interpretation as in Figure 4. The median performance of `E-ON-MMU/R-MMU` across 50 random utility vectors $\boldsymbol{u}^\star$ is shown with red dots. The median performance of `E-RAND/R-MMU` is shown with blue crosses. The median performance of `E-AC/R-AC` (resp. `E-AC/R-MMU` is shown with green crosses (resp. triangles).

information to maximize utility can be gained (in the worst-case) by asking queries. Thus, we focus our analysis on the min-max regret case.

Offline Preference Elicitation: Optimality and Scalability Results. We compare our offline min-max regret elicitation approach `E-OFF-MMR/R-MMR` to random elicitation `E-RAND/R-MMR`. The results are summarized on Figure 10. From the figure, it can be see that `E-OFF-MMR/R-MMR` outperforms `E-RAND/R-MMR`. Indeed, the standardized worst-case regret using our elicitation approach drops to by 20 percentage points after just one query and by 70 percentage points after 10 queries. On the other hand, the standardized worst-case regret remains equal to 1 after one random query and does not drop below 0.85 after even 10 random queries (it never decreases after 2 random queries). We also note that the evaluation time for a random set of queries and the time to compute a near optimal set of queries are comparable.

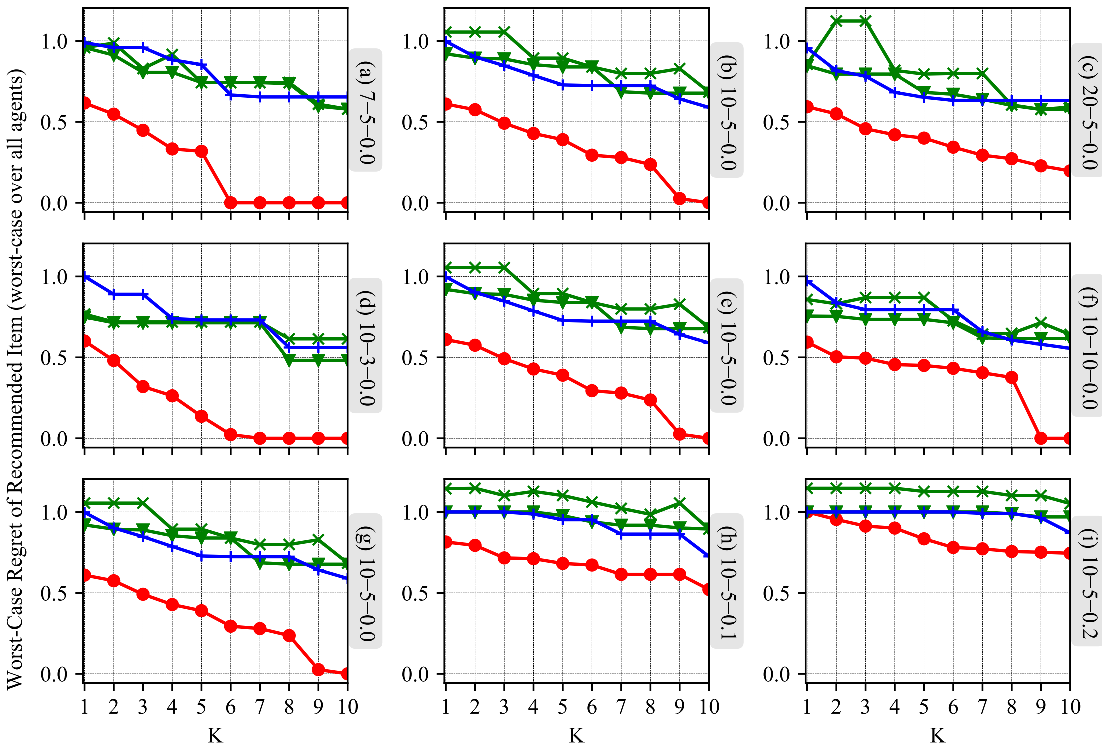

Figure 9 Optimality results for the online regret averse active preference elicitation problem ($\mathcal{P}^K_{\text{on,regret}}$) on synthetic data. The facet labels have the same interpretation as in Figure 6. The median performance of `E-ON-MMR/R-MMR` across 50 random utility vectors $\boldsymbol{u}^\star$ is shown with red dots. The median performance of `E-RAND/R-MMR` is shown with blue crosses. The median performance of `E-AC/R-AC` (resp. `E-AC/R-MMR` is shown with green crosses (resp. triangles).

Online Preference Elicitation: Optimality Results. For our last set of experiments, we evaluate the performance of our online regret averse preference elicitation approach `E-ON-MMR/R-MMR` on the real dataset from Section 8.2. We compare it to existing techniques from the literature, see Section 8.3. The results are summarized on Figure 11. From the figure, it can be seen that the worst-case normalized regret of our proposed approach drops by almost 80 percentage points after just 4 queries, and by 90 percentage points after 10 queries. On the other hand, the best approach from the literature has a worst-case normalized regret higher than 0.78 after even 10 queries. In addition, our approach is able to recommend a top ranked item after only seven queries whereas the best approach from the literature offered the 6th ranked item after 10 queries.

Table 1 Comparison between max-min utility and min-max regret based queries on a synthetic dataset with $I = 10$ items and $J = 10$ features ($\Gamma = 0$). The max-min utility queries $\boldsymbol{\iota}^\star_{\rm u}$ and min-max regret queries $\boldsymbol{\iota}^\star_{\rm r}$ are computed by solving the MBLP reformulations of Problems ($\mathcal{P}^K_{\rm off,risk}$) and ($\mathcal{P}^K_{\rm off,regret}$), respectively, using the column-and-constraint generation algorithm. The decrease in normalized worst-case utility refers to the drop in normalized worst-case utility experienced by employing the min-max regret rather than max-min utility solution, computed as $(u_{\rm nwc}(\boldsymbol{\iota}^\star_{\rm u}) - u_{\rm nwc}(\boldsymbol{\iota}^\star_{\rm r}))$. Similarly, the decrease in normalized worst-case regret refers to the drop in normalized worst-case regret experienced by employing the min-max regret rather than max-min utility solution, computed as $(r_{\rm nwc}(\boldsymbol{\iota}^\star_{\rm u}) - r_{\rm nwc}(\boldsymbol{\iota}^\star_{\rm r}))$. All approaches were given a one hour time limit.

	Max-Min Utility Solution		Min-Max Regret Solution			
K	Normalized Worst-Case Utility	Normalized Worst-Case Regret	Normalized Worst-Case Utility	Normalized Worst-Case Regret	Decrease in Normalized Worst-Case Utility (p.p.)	Decrease in Normalized Worst-Case Regret (p.p.)
2	0.780	0.642	0.689	0.594	13.8 [**NOTE DCM:** *should be 9.2?*]	4.8
4	0.783	0.620	0.691	0.532	9.2	8.8
6	0.833	0.566	0.712	0.522	12.1	4.4
8	0.879	0.550	0.816	0.499	6.3	5 [**NOTE DCM:** *sho*
10	0.912	0.510	0.844	0.495	6.8	1.5

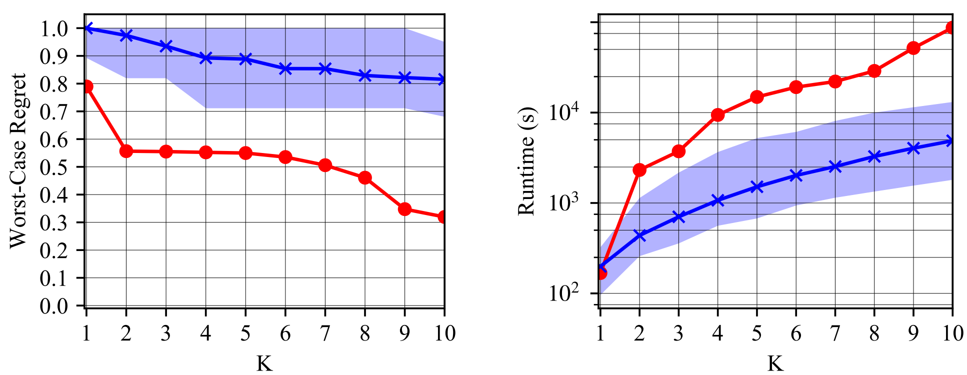

Figure 10 Optimality (left) and scalability (right) results for the offline regret averse active preference elicitation problem ($\mathcal{P}^K_{\rm off,regret}$) on the real dataset. Approach `E-OFF-MMR` is shown with red dots and the median (and range of) performance of `E-RAND` across 13 random queries is shown with blue crosses (shade).

9. Conclusion

In this paper, we proposed novel formulations for the offline and online risk averse and regret averse active preference elicitation problems with and without inconsistencies. These take the form of robust optimization problems with decision-dependent information discovery. We studied the complexity of these problems and

Table 2 Evaluation results of symmetry breaking constraints, CCG algorithm, and CCG-based greedy heuristic approach on three synthetic datasets with $I \in \{5, 10, 20\}$ items and $J = 10$ features ($\Gamma = 0$); A dash indicates that the optimal solution was not found within the allotted 3 hour time limit.

I	K	MBLP		MBLP + Symm. Break.		CCG + Symm. Break.		Heuristic	
		Normalized Objective Value	Solve Time (sec)	Normalized Objective Value	Solve Time (sec)	Normalized Objective Value	Solve Time (sec)	Normalized Objective Value	Solve Time (sec)
5	2	0.800	0.34	0.800	0.25	0.800	0.77	0.800	0.51
5	4	0.932	622.64	0.932	2.45	0.932	5.48	0.932	1.40
5	6	0.978	6094.79	0.978	76.21	0.978	10.80	0.932	2.25
5	8	–	–	0.978	1147.51	0.978	40.91	0.972	3.53
5	10	–	–	1.000	5082.32	1.000	0.86	1.000	4.24
10	2	0.780	2.66	0.780	2.61	0.780	5.54	0.780	2.13
10	4	–	–	0.879	1141.58	0.879	6054.11	0.783	5.67
10	6	–	–	–	–	–	–	0.812	12.53
10	8	–	–	–	–	–	–	0.879	21.44
10	10	–	–	–	–	–	–	0.879	31.79
20	2	0.646	334.33	0.646	202.75	0.646	338.17	0.646	5.21
20	4	–	–	–	–	–	–	0.658	23.24
20	6	–	–	–	–	–	–	0.765	51.16
20	8	–	–	–	–	–	–	0.786	84.35
20	10	–	–	–	–	–	–	0.812	162.32

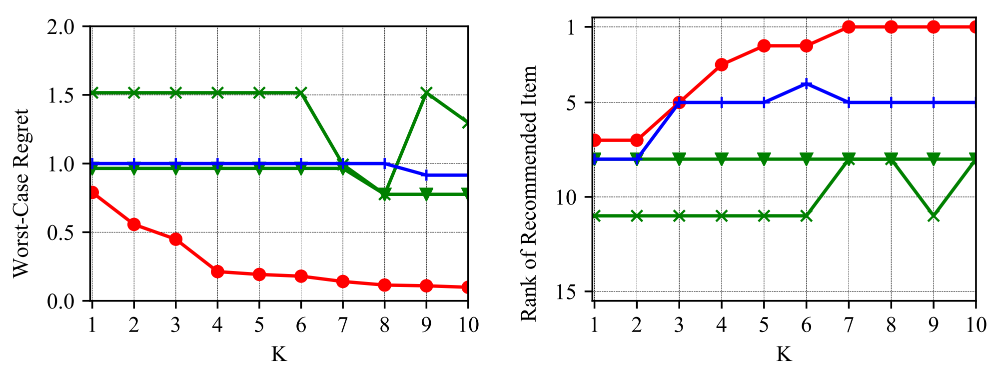

Figure 11 Worst-case regret (left) and rank (right) of the recommended item for the online regret averse active preference elicitation problem ($\mathcal{P}^K_{\text{on,regret}}$) on the real dataset. The median performance of `E-ON-MMR/R-MMR` across 13 random utility vectors $\boldsymbol{u}^\star$ is shown with red dots. The median performance of `E-RAND/R-MMR` is shown with blue crosses. The median performance of `E-AC/R-AC` (resp. `E-AC/R-MMR` is shown with green crosses (resp. triangles).

provided exact reformulations and conservative approximations of the offline and online problems, respectively. We provided efficient solutions procedures and performed extensive computational experiments that showed the superiority of our approach on both synthetic data and real data from the homeless management information system. In the future, we plan to deploy this algorithm on Amazon Mechanical Turk[10] and on policy-makers at the Los Angeles Homeless Services Authority to be able to elicit their preferences and recommend housing allocation policies that best meet their needs to help mitigate homelessness.

Notes

[1]See `https://www.google.com/maps`.

[2]See `https://www.apple.com/ios/maps/`.

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

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

[5]`https://controllerdata.lacity.org/`

[6]See `https://www.uber.com`.

[7]See `https://www.lyft.com`.

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

[9]TAY VI-SPDAT stands for Transition Age Youth Vulnerability Index-Service Prioritization Decision Assistance Tool; it is the tool used to assess homeless youth when they enter the housing allocation system.

[10]`https://www.mturk.com`

Acknowledgments

The authors gratefully acknowledge support from Schmidt Futures, from the James H. Zumberge Faculty Research and Innovation Fund at the University of Southern California, and from the U.S. Department of Transportation METRANS University Transportation Center and National Center for Sustainable Transportation under Award No. DTRT13-G-UTC57. The authors are grateful to Erin Cox, Policy Manager at LAHSA, and Maggie Potthoff, Policy Supervisor at LAHSA, for their support of this work and for the very interesting discussions that helped shape this research. They would also like to thank Andrés Gómez for vibrant conversations on polyhedral theory.

References

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, ISSN 00129682, URL `http://dx.doi.org/10.2307/1907921`.

Arrow KJ (1963) *Social choice and individual values* (Yale University Press), 2nd edition.

Azizi MJ, Vayanos P, Wilder B, Rice E, Tambe M (2018) Designing Fair, Efficient, and Interpretable Policies for Prioritizing Homeless Youth for Housing Resources. *Proceedings of the 15th International Conference on the Integration of Constraint Programming, Artificial Intelligence, and Operations Research.*

Bandi C, Bertsimas D (2012) Tractable stochastic analysis in high dimensions via robust optimization. *Mathematical programming* 1–48.

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.

Basciftci B, Ahmed S, Shen S (2019) Distributionally Robust Facility Location Problem under Decision-dependent Stochastic Demand.

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

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

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, ISSN 1573-0565, URL `http://dx.doi.org/10.1007/s10994-017-5633-9`.

Bertsimas D, Dunning I (2016) Multistage robust mixed-integer optimization with adaptive partitions. *Operations Research* 64(4):980–998, URL `http://dx.doi.org/10.1287/opre.2016.1515`.

Bertsimas D, Farias VF, Trichakis N (2013) Fairness, efficiency, and flexibility in organ allocation for kidney transplantation. *Operations Research* 61(1):73–87.

Bertsimas D, Georghiou A (2015) Design of near optimal decision rules in multistage adaptive mixed-integer optimization. *Operations Research* 63(3):610–627, URL `http://dx.doi.org/10.1287/opre.2015.1365`.

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, URL `http://dx.doi.org/10.1287/opre.2013.1209`.

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

Bodur M, Luedtke JR (2018) Two-stage linear decision rules for multi-stage stochastic programming. *Mathematical Programming* ISSN 1436-4646, URL `http://dx.doi.org/10.1007/s10107-018-1339-4`.

Boutilier C (2002) A POMDP formulation of preference elicitation problems. *Proceedings of the Eighteenth AAAI Conference on Artificial Intelligence (AAAI2002)*, 239–246.

Boutilier C, Patrascu R, Poupart P, Schuurmans D (2006) Constraint-based optimization and utility elicitation using the minimax decision criterion. *Artificial Intelligence* 170(8-9):686–713.

Braziunas D, Boutilier C (2010) Assessing regret-based preference elicitation with the UTPREF recommendation system. *Proceedings of the 11th ACM conference on Electronic commerce (EC2010)*, 219–228 (ACM).

Brochu E, de Freitas N, Ghosh A (2007) Active preference learning with discrete choice data. *Proceedings of the 20th International Conference on Neural Information Processing Systems (NIPS'07)*, 409–416.

Carroll JD, Green PE (1997) Psychometric Methods in Marketing Research: Part II, Multidimensional Scaling. *Journal of Marketing Research* 34(2):193, ISSN 00222437, URL http://dx.doi.org/10.2307/3151858.

Chajewska U, Koller D, Parr R (2000) Making rational decisions using adaptive utility elicitation. *Proceedings of the Seventeenth Conference on Artificial Intelligence (AAAI2000)*, 363–369.

Chan H, Rice E, Vayanos P, Tambe M, Morton M (2017) Evidence From the Past: AI Decision Aids to Improve Housing Systems for Homeless Youth. *AAAI 2017 Fall Symposium Series*.

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, ISSN 0377-2217, URL http://dx.doi.org/https://doi.org/10.1016/j.ejor.2019.05.045.

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.

Debreu G (1954) *Decision Process* (New York: John Wiley), URL http://www.economics-ejournal.org/economics/journalarticles/2009-12/references/Debreu1954.

Domshlak C, Hüllermeier E, Kaci S, Prade H (2011) Preferences in AI: An overview. *Artificial Intelligence* 175(7):1037–1052, ISSN 0004-3702, URL http://dx.doi.org/https://doi.org/10.1016/j.artint.2011.03.004.

Dyer JS, Fishburn PC, Steuer RE, Wallenius J, Zionts S (1992) Multiple Criteria Decision Making, Multiattribute Utility Theory: The Next Ten Years. *Management Science* 38(5):645–654, ISSN 0025-1909, URL http://dx.doi.org/10.1287/mnsc.38.5.645.

Dyer JS, Sarin RK (1979) Measurable Multiattribute Value Functions. *Operations Research* 27(4):810–822, URL http://dx.doi.org/10.1287/opre.27.4.810.

Garey MR, Johnson DS (1979) Computers and intractability: a guide to the theory of NP-completeness. *San Francisco, LA: Freeman* 58.

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.

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

Green PE, Srinivasan V (1990) Conjoint Analysis in Marketing: New Developments with Implications for Research and Practice. *Journal of Marketing* 54(4):3, ISSN 00222429, URL `http://dx.doi.org/10.2307/1251756`.

Gupta V, Grossmann IE (2011) Solution strategies for multistage stochastic programming with endogenous uncertainties. *Computers & Chemical Engineering* 35(11):2235–2247, ISSN 0098-1354, URL `http://dx.doi.org/https://doi.org/10.1016/j.compchemeng.2010.11.013`.

Hanasusanto GA, Kuhn D, Wiesemann W (2015) K-Adaptability in two-stage robust binary programming. *Operations Research* 63(4):877–891, URL `http://dx.doi.org/10.1287/opre.2015.1392`.

Hillier FS (2012) *Introduction To Operations Research* (McGraw-Hill Education (India) Pvt Limited), ISBN 9780071333467, URL `https://books.google.com/books?id=fsCjAgAAQBAJ`.

Johnson RM (1987) Accuracy of Utility Estimation in ACA. *Sawtooth Software* .

Johnson RM (1991) Comment on "Adaptive Conjoint Analysis: Some Caveats and Suggestions". *Journal of Marketing Research* ISSN 00222437, URL `http://dx.doi.org/10.2307/3172810`.

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, URL `http://dx.doi.org/10.1257/jep.5.1.193`.

Kahneman D, Tversky A (1979) Prospect Theory: an analysis of decision under risk. *Econometrica* 47(2):263, URL http://dx.doi.org/10.2307/1914185.

Keeney RL, L KR, Raiffa H (1993) *Decisions with Multiple Objectives: Preferences and Value Trade-Offs*. Wiley series in probability and mathematical statistics. Applied probability and statistics (Cambridge University Press), ISBN 9780521438834, URL https://books.google.com/books?id=GPE6ZAqGrnoC.

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 (2018) Robust optimization for decision-making under endogenous uncertainty. *Computers and Chemical Engineering* ISSN 00981354, URL http://dx.doi.org/10.1016/j.compchemeng.2018.01.006.

Luo F, Mehrotra S (2019) Service Center Location Problem with Decision Dependent Utilities.

Milnor J (1954) Games against nature. Thrall R, Coombs C, Davis R, eds., *Decision Processes* (Wiley, New York).

Nemhauser G, Wolsey L (1988) *Integer and Combinatorial Optimization*. URL http://dx.doi.org/10.1002/9781118627372.

Neumann JV, Morgenstern O (1944) *Theory of Games and Economic Behavior* (University Press Group Ltd), ISBN 0691130612.

Nohadani O, Roy A (2017) Robust optimization with time-dependent uncertainty in radiation therapy. *IISE Transactions on Healthcare Systems Engineering* 7(2):81–92, URL http://dx.doi.org/10.1080/24725579.2017.1296907.

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

Noyan N, Rudolf G, Lejeune M (2018) Distributionally Robust Optimization with Decision-Dependent Ambiguity.

Popper B (2017) Google announces over 2 billion monthly active devices on Android.

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), ISSN 15265528, URL `http://dx.doi.org/10.1287/ijoc.2016.0696`.

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

Ryu M, Jiang R (2019) Nurse Staffing under Absenteeism: A Distributionally Robust Optimization Approach.

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

Simon H (1955) A Behavioral Model of Rational Choice. *The Quarterly Journal of Economics* 69(1):99–118, URL `http://econpapers.repec.org/RePEc:oup:qjecon:v:69:y:1955:i:1:p:99-118`.

Slovic P, Fischhoff B, Lichtenstein S (1977) Behavioral Decision Theory. *Annual Review of Psychology* ISSN 0066-4308, URL `http://dx.doi.org/10.1146/annurev.ps.28.020177.000245`.

Spacey SA, Wiesemann W, Kuhn D, Luk W (2012) Robust software partitioning with multiple instantiation. *INFORMS Journal on Computing* 24(3):500–515, URL `http://dx.doi.org/10.1287/ijoc.1110.0467`.

Stoye J (2011) Statistical decisions under ambiguity. *Theory and Decision* 70(2):129–148, ISSN 0040-5833, URL `http://dx.doi.org/10.1007/s11238-010-9227-2`.

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

Toubia O, Hauser J, Garcia R (2007) Probabilistic Polyhedral Methods for Adaptive Choice-Based Conjoint Analysis: Theory and Application. *Marketing Science* 26(5):596–610, ISSN 0732-2399, URL `http://dx.doi.org/10.1287/mksc.1060.0257`.

Toubia O, Hauser JR, Simester DI (2004) Polyhedral Methods for Adaptive Choice-Based Conjoint Analysis. *Journal of Marketing Research* 41(1):116–131, ISSN 0022-2437, URL `http://dx.doi.org/10.1509/jmkr.41.1.116.25082`.

Toubia O, Simester DI, Hauser JR, Dahan E (2003) Fast Polyhedral Adaptive Conjoint Estimation. *Marketing Science* 22(3):273–303, URL http://dx.doi.org/10.1287/mksc.22.3.273.17743.

Vayanos P, Georghiou A, Yu H (2019) Robust optimization with decision-dependent information discovery. Under review at Management Science.

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.

Von Neumann J (1928) Zur Theorie der Gesellschaftsspiele. *Mathematische Annalen* ISSN 00255831, URL http://dx.doi.org/10.1007/BF01448847.

Wang T, Boutilier C (2003) Incremental utility elicitation with the minimax regret decision criterion. *Proceedings of the 18th international joint conference on Artificial intelligence (IJCAI'03)*, volume 3, 309–316.

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.

Zhang X, Kamgarpour M, Georghiou A, Goulart P, Lygeros J (2017) Robust optimal control with adjustable uncertainty sets. *Automatica* 75(Supplement C):249–259, ISSN 0005-1098, URL http://dx.doi.org/https://doi.org/10.1016/j.automatica.2016.09.016.

Zhao Z, Li H, Wang J, Kephart J, Mattei N, Su H, Xia L (2018) A cost-effective framework for preference elicitation and aggregation. *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI'18)*.

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

Zionts S, Wallenius J (1976) An Interactive Programming Method for Solving the Multiple Criteria Problem. *Management Science* 22(6):652–663, ISSN 0025-1909, URL http://dx.doi.org/10.1287/mnsc.22.6.652.

E-Companion

EC.1. Proofs of Statements in Sections 3 and 4

The following proposition is needed in the proof of Lemma 1. It shows that we can eliminate the dependence of $\mathcal{S}(\boldsymbol{\iota})$ on $\boldsymbol{\iota}$ in Problem ($\mathcal{P}^K_{\text{off,risk}}$).

Proposition EC.1. *Problem* ($\mathcal{P}^K_{\text{off,risk}}$) *is equivalent to*

$$\underset{\boldsymbol{\iota}\in\mathcal{C}^K}{\text{maximize}} \quad \min_{\boldsymbol{s}\in\mathcal{S}^K} \quad \max_{\boldsymbol{x}\in\mathcal{R}} \quad \min_{\boldsymbol{u}\in\mathcal{U}(\boldsymbol{\iota},\boldsymbol{s})} \quad \boldsymbol{u}^\top\boldsymbol{x} \tag{EC.1}$$

in the sense that the two problems have the same optimal objective value and the same sets of optimal solutions.

Proof Fix $\boldsymbol{\iota}\in\mathcal{C}^K$. Fix $\boldsymbol{s}\in\mathcal{S}^K\backslash\{\mathcal{S}(\boldsymbol{\iota})\}$. Then, $\mathcal{U}(\boldsymbol{\iota},\boldsymbol{s})=\emptyset$. Since the choice of $\boldsymbol{s}\in\mathcal{S}^K\backslash\{\mathcal{S}(\boldsymbol{\iota})\}$ was arbitrary, it follows from $\mathcal{R}\neq\emptyset$ that

$$\min_{\boldsymbol{s}\in\mathcal{S}^K\backslash\{\mathcal{S}(\boldsymbol{\iota})\}} \quad \max_{\boldsymbol{x}\in\mathcal{R}} \quad \min_{\boldsymbol{u}\in\mathcal{U}(\boldsymbol{\iota},\boldsymbol{s})} \quad \boldsymbol{u}^\top\boldsymbol{x} \;=\; +\infty.$$

Therefore,

$$\min_{\boldsymbol{s}\in\mathcal{S}^K\backslash\{\mathcal{S}(\boldsymbol{\iota})\}} \quad \max_{\boldsymbol{x}\in\mathcal{R}} \quad \min_{\boldsymbol{u}\in\mathcal{U}(\boldsymbol{\iota},\boldsymbol{s})} \quad \boldsymbol{u}^\top\boldsymbol{x} \quad \geq \quad \min_{\boldsymbol{s}\in\mathcal{S}(\boldsymbol{\iota})} \quad \max_{\boldsymbol{x}\in\mathcal{R}} \quad \min_{\boldsymbol{u}\in\mathcal{U}(\boldsymbol{\iota},\boldsymbol{s})} \quad \boldsymbol{u}^\top\boldsymbol{x}.$$

This implies that

$$\min_{\boldsymbol{s}\in\{\mathcal{S}(\boldsymbol{\iota})\}\cup\mathcal{S}^K\backslash\{\mathcal{S}(\boldsymbol{\iota})\}} \quad \max_{\boldsymbol{x}\in\mathcal{R}} \quad \min_{\boldsymbol{u}\in\mathcal{U}(\boldsymbol{\iota},\boldsymbol{s})} \quad \boldsymbol{u}^\top\boldsymbol{x} \quad = \quad \min_{\boldsymbol{s}\in\mathcal{S}(\boldsymbol{\iota})} \quad \max_{\boldsymbol{x}\in\mathcal{R}} \quad \min_{\boldsymbol{u}\in\mathcal{U}(\boldsymbol{\iota},\boldsymbol{s})} \quad \boldsymbol{u}^\top\boldsymbol{x}.$$

Since the choice of $\boldsymbol{\iota}\in\mathcal{C}^K$ was arbitrary, this concludes the proof. $\square$

Proof of Lemma 1 From Proposition EC.1, it suffices to show that Problems (EC.1) and (2) are equivalent. To this end, fix any $\boldsymbol{\iota}\in\mathcal{C}^K$. We show that

$$\min_{\boldsymbol{s}\in\mathcal{S}^K} \quad \max_{\boldsymbol{x}\in\mathcal{R}} \quad \min_{\boldsymbol{u}\in\mathcal{U}(\boldsymbol{\iota},\boldsymbol{s})} \quad \boldsymbol{u}^\top\boldsymbol{x} \tag{EC.2}$$

and

$$\min_{\boldsymbol{s}\in\mathcal{S}^K} \quad \max_{\boldsymbol{x}\in\mathcal{R}} \quad \min_{\boldsymbol{u}\in\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s})} \quad \boldsymbol{u}^\top\boldsymbol{x} \tag{EC.3}$$

attain the same optimal value.

Since $\mathcal{U}(\boldsymbol{\iota},\boldsymbol{s})\subseteq\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s})$, the optimal objective value of Problem (EC.3) constitutes a lower bound on the optimal objective value of Problem (EC.2).

For the converse part, observe that, by virtue of the linearity of $\boldsymbol{u}^\top \boldsymbol{x}$ and the boundedness of $\mathcal{U}(\boldsymbol{\iota}, \boldsymbol{s})$, Problems (EC.2) and (EC.3) have the same objective value for all $\boldsymbol{s} \in \mathcal{S}^K$ such that $\mathcal{U}(\boldsymbol{\iota}, \boldsymbol{s}) \neq \emptyset$ since, in such cases, $\widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s}) = \mathrm{cl}(\mathcal{U}(\boldsymbol{\iota}, \boldsymbol{s}))$. Let $\boldsymbol{s}^\star$ be optimal in Problem (EC.3). If $\mathcal{U}(\boldsymbol{\iota}, \boldsymbol{s}^\star) \neq \emptyset$, the proof is complete. Suppose instead that $\mathcal{U}(\boldsymbol{\iota}, \boldsymbol{s}^\star) = \emptyset$. Since $\mathcal{U}(\boldsymbol{\iota}, \boldsymbol{s}^\star) = \emptyset$ while $\widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s}^\star) \neq \emptyset$, this implies that there exist implied equalities in the set $\widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s}^\star)$. In particular, since $\mathcal{U}^0$ is full-dimensional, see Assumption 2, the implied equalities are all associated with response constraints. Let us collect the indices κ of all those implied equalities in the set $\mathcal{E}$, i.e., let

$$\mathcal{E} := \left\{ \kappa \in \mathcal{K} : \boldsymbol{s}_\kappa \in \{-1, 1\}, \, \boldsymbol{u}^\top (\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}) = 0 \quad \forall \boldsymbol{u} \in \widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s}^\star) \right\}.$$

Define

$$\boldsymbol{s}'_\kappa := \begin{cases} 1 & \text{if } \boldsymbol{s}_\kappa = 1 \text{ and } \kappa \notin \mathcal{E}, \\ 0 & \text{if } \boldsymbol{s}_\kappa = 0 \text{ or } \kappa \in \mathcal{E}, \\ -1 & \text{if } \boldsymbol{s}_\kappa = -1 \text{ and } \kappa \notin \mathcal{E}. \end{cases}$$

This definition of $\boldsymbol{s}'$ ensures that equality constraints that are implicit in the set $\widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s}^\star)$ are made explicit in the set $\widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s}')$. By construction, $\widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s}') = \widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s}^\star)$ and thus $\boldsymbol{s}'$ is optimal in Problem (EC.3). Moreover, since $\widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s}')$ is non-empty, it follows by Proposition 2.3 in Nemhauser and Wolsey (1988) that $\mathcal{U}(\boldsymbol{\iota}, \boldsymbol{s}')$ is non-empty. Thus, $\widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s}') = \mathrm{cl}(\mathcal{U}(\boldsymbol{\iota}, \boldsymbol{s}'))$, implying that the objective value attained by $\boldsymbol{s}'$ in Problem (EC.2) equals the optimal objective value of Problem (EC.3). This in turn yields the required result that the optimal objective value of Problem (EC.3) constitutes an upper bound on the optimal objective value of Problem (EC.2).

Thus, the optimal values of Problems (EC.2) and (EC.3) are equal, which concludes the proof. □

Proof of Observation 1 Fix any $\boldsymbol{\iota} \in \mathcal{C}^K$. Let $\boldsymbol{s}^\star$ be optimal in

$$\min_{\boldsymbol{s} \in \mathcal{S}^K} \quad \max_{\boldsymbol{x} \in \mathcal{R}} \quad \min_{\boldsymbol{u} \in \widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s})} \quad \boldsymbol{u}^\top \boldsymbol{x} \tag{EC.4}$$

and define $\boldsymbol{s}' \in \widetilde{\mathcal{S}}^K \subset \mathcal{S}^K$ through

$$\boldsymbol{s}'_\kappa := \begin{cases} \boldsymbol{s}^\star_\kappa & \text{if } \boldsymbol{s}^\star_\kappa \neq 0 \\ 1 & \text{else.} \end{cases}$$

Then, $\widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s}') \supseteq \widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s}^\star)$, implying that the objective value attained by $\boldsymbol{s}'$ in Problem (EC.4) is no greater than the objective value attained by $\boldsymbol{s}^\star$. Thus, $\boldsymbol{s}'$ is also optimal in (EC.4). Since $\boldsymbol{s}' \in \widetilde{\mathcal{S}}^K$, this implies that we can always restrict our search for an optimal solution to Problem (EC.4) to the set $\widetilde{\mathcal{S}}^K$. As the choice of $\boldsymbol{\iota} \in \mathcal{C}^K$ was arbitrary, this concludes the proof. $\square$

Proof of Theorem 1 *(a)* Fix any $K \in \mathbb{N}$, $\boldsymbol{\iota} \in \mathcal{C}^K$, and $\boldsymbol{s} \in \widetilde{\mathcal{S}}^K$ such that $\widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s}) \neq \emptyset$. Then, $\widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s})$ is convex. If $\mathcal{R}$ is convex, it readily follows from the minimax theorem (see Von Neumann (1928)) that

$$\min_{\boldsymbol{u} \in \widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s})} \max_{\boldsymbol{x} \in \mathcal{R}} \boldsymbol{u}^\top \boldsymbol{x} = \max_{\boldsymbol{x} \in \mathcal{R}} \min_{\boldsymbol{u} \in \widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s})} \boldsymbol{u}^\top \boldsymbol{x}.$$

Since the choices of $K \in \mathbb{N}$, $\boldsymbol{\iota} \in \mathcal{C}^K$, and $\boldsymbol{s} \in \widetilde{\mathcal{S}}^K$ such that $\widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s}) \neq \emptyset$ were arbitrary, it holds that (under the premise of the theorem statement) Problem ($\widetilde{\mathcal{P}}^K_{\text{off,risk}}$) is equivalent to

$$\underset{\boldsymbol{\iota} \in \mathcal{C}^K}{\text{maximize}} \min_{\boldsymbol{s} \in \widetilde{\mathcal{S}}^K} \min_{\boldsymbol{u} \in \widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s})} \max_{\boldsymbol{x} \in \mathcal{R}} \boldsymbol{u}^\top \boldsymbol{x}. \tag{EC.5}$$

Fix $K \in \mathbb{N}$. Then,

$$\begin{aligned} & \max_{\boldsymbol{\iota} \in \mathcal{C}^K} \min_{\boldsymbol{s} \in \widetilde{\mathcal{S}}^K} \min_{\boldsymbol{u} \in \widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s})} \max_{\boldsymbol{x} \in \mathcal{R}} \boldsymbol{u}^\top \boldsymbol{x} \\ = & \max_{\boldsymbol{\iota} \in \mathcal{C}^K} \min_{\boldsymbol{u} \in \bigcup_{\boldsymbol{s} \in \mathcal{S}^K} \widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s})} \max_{\boldsymbol{x} \in \mathcal{R}} \boldsymbol{u}^\top \boldsymbol{x} \\ = & \max_{\boldsymbol{\iota} \in \mathcal{C}^K} \min_{\boldsymbol{u} \in \mathcal{U}^0} \max_{\boldsymbol{x} \in \mathcal{R}} \boldsymbol{u}^\top \boldsymbol{x} \\ = & \min_{\boldsymbol{u} \in \mathcal{U}^0} \max_{\boldsymbol{x} \in \mathcal{R}} \boldsymbol{u}^\top \boldsymbol{x} \\ = & \max_{\boldsymbol{x} \in \mathcal{R}} \min_{\boldsymbol{u} \in \mathcal{U}^0} \boldsymbol{u}^\top \boldsymbol{x}, \end{aligned} \tag{EC.6}$$

where the first equality follows since, for any $\boldsymbol{\iota} \in \mathcal{C}^K$, we have

$$\left\{ \boldsymbol{u} \in \widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s}) : \boldsymbol{s} \in \widetilde{\mathcal{S}}^K \right\} = \bigcup_{\boldsymbol{s} \in \widetilde{\mathcal{S}}^K} \widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s}),$$

and the second equality follows by definitions of $\widetilde{\mathcal{S}}^K$ and $\widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s})$. If the recommendation set $\mathcal{R}$ is a polyhedron, then

$$\max_{\boldsymbol{x} \in \mathcal{R}} \min_{\boldsymbol{u} \in \mathcal{U}^0} \boldsymbol{u}^\top \boldsymbol{x}$$

can be reformulated equivalently as a linear program of size polynomial in the size of the input parameters, see e.g., Ben-Tal et al. (2009), and is thus polynomially solvable. Since the choice of $K \in \mathbb{N}$ above was arbitrary, this concludes the first proof of the first item.

(b) We use a reduction from the following decision problem that is known to be $\mathcal{NP}$-complete, see Garey and Johnson (1979).

PARTITION.

Instance. Given a set $\mathcal{A}$ of elements $\mathcal{A} := \{1, \ldots, n\}$ with associated positive integer weights $\boldsymbol{w}_i \in \mathbb{N}_+$, $i \in \mathcal{A}$, such that $\sum_{i \in \mathcal{A}} \boldsymbol{w}_i = 2W$.

Question. Does there exist a partition of $\mathcal{A}$ into two subsets, $\mathcal{X}$ and $\mathcal{A} \backslash \mathcal{X}$, such that $\sum_{i \in \mathcal{X}} \boldsymbol{w}_i = \sum_{i \in \mathcal{A} \backslash \mathcal{X}} \boldsymbol{w}_i = W$?

We aim to reduce the partition problem to evaluating the objective function of an instance of Problem ($\widetilde{\mathcal{P}}^K_{\text{off,risk}}$). To this end, fix an instance $(n, \boldsymbol{w}, W)$ of the partition problem. Set $J := n+1$, $K = n$, and $\mathcal{U}^0 := \{\boldsymbol{u} \in [0,1]^J \,:\, \boldsymbol{u}_J = 0.5\}$. Also, for each $\kappa \in \mathcal{K}$, let $\boldsymbol{\iota}_1^\kappa := \mathbf{e}_\kappa$ and $\boldsymbol{\iota}_2^\kappa := \mathbf{e}_J$. Then, given a choice $\boldsymbol{s} \in \widetilde{\mathcal{S}}^K$, we have

$$\widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s}) := \left\{ \begin{array}{lll} \boldsymbol{u} \in [0,1]^J \,:\, \boldsymbol{u}_J = 0.5 & & \\ \qquad\qquad\quad \boldsymbol{u}_\kappa \geq 0.5 & \forall \kappa \in \mathcal{K} \,:\, \boldsymbol{s}_\kappa = 1 \\ \qquad\qquad\quad \boldsymbol{u}_\kappa \leq 0.5 & \forall \kappa \in \mathcal{K} \,:\, \boldsymbol{s}_\kappa = -1 \end{array} \right\}.$$

Finally, we define the recommendation set $\mathcal{R} \subseteq \mathbb{R}^J$ through $\mathcal{R} := \{\boldsymbol{x}^1, \boldsymbol{x}^2\}$, where $\boldsymbol{x}^1 := (2\boldsymbol{w}^\top, -2W)^\top$ and $\boldsymbol{x}^2 := ((-2\boldsymbol{w})^\top, 6W)^\top$. For any fixed $\boldsymbol{s} \in \widetilde{\mathcal{S}}^K$, define

$$Z(\boldsymbol{s}) \quad := \quad \max_{\boldsymbol{x} \in \mathcal{R}} \quad \min_{\boldsymbol{u} \in \widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s})} \quad \boldsymbol{u}^\top \boldsymbol{x}.$$

Then, we have

$$\begin{aligned} Z(\boldsymbol{s}) \quad &= \quad \max \left\{ \left(\min_{\boldsymbol{u} \in \widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s})} \sum_{i \in \mathcal{A}} 2\boldsymbol{u}_i \boldsymbol{w}_i - 2W\boldsymbol{u}_J \right), \left(\min_{\boldsymbol{u} \in \widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s})} \sum_{i \in \mathcal{A}} -2\boldsymbol{u}_i \boldsymbol{w}_i + 6W\boldsymbol{u}_J \right) \right\} \\ &= \quad \max \left\{ \left(\min_{\boldsymbol{u} \in \widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s})} \sum_{i \in \mathcal{A}} 2\boldsymbol{u}_i \boldsymbol{w}_i \right) - W, \left(\min_{\boldsymbol{u} \in \widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s})} \sum_{i \in \mathcal{A}} -2\boldsymbol{u}_i \boldsymbol{w}_i + 2W \right) + W \right\} \\ &= \quad \max \left\{ \left(\min_{\boldsymbol{u} \in \widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s})} \sum_{i \in \mathcal{A}} 2\boldsymbol{u}_i \boldsymbol{w}_i \right) - W, \left(\min_{\boldsymbol{u} \in \widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s})} \sum_{i \in \mathcal{A}} -2\boldsymbol{u}_i \boldsymbol{w}_i + \sum_{i \in \mathcal{A}} \boldsymbol{w}_i \right) + W \right\} \\ &= \quad \max \left\{ \left(\min_{\boldsymbol{u} \in \widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s})} \sum_{i \in \mathcal{A}} 2\boldsymbol{u}_i \boldsymbol{w}_i \right) - W, \left(\min_{\boldsymbol{u} \in \widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s})} \sum_{i \in \mathcal{A}} -2(\boldsymbol{u}_i - 0.5) \boldsymbol{w}_i \right) + W \right\} \\ &= \quad \max \left\{ \sum_{\kappa \in \mathcal{K} : \boldsymbol{s}_\kappa = 1} \boldsymbol{w}_\kappa - W, \sum_{\kappa \in \mathcal{K} : \boldsymbol{s}_\kappa = 1} -\boldsymbol{w}_\kappa + W \right\} \\ &\geq \quad 0. \end{aligned}$$

Now, we claim that we are given a "yes" instance of PARTITION if and only the objective value of $\boldsymbol{\iota}$ in the constructed instance of Problem ($\widetilde{\mathcal{P}}^K_{\text{off,risk}}$) is 0. To this end, note that the objective value of $\boldsymbol{\iota}$ is given by

$$\underset{\boldsymbol{s} \in \widetilde{\mathcal{S}}^K}{\text{minimize}} \quad Z(\boldsymbol{s}) \tag{EC.7}$$

and is lower bounded by 0. If there exists $\mathcal{X} \subset \mathcal{A}$ such that $\sum_{i\in\mathcal{X}} \boldsymbol{w}_i = W$, then the solution $\boldsymbol{s} \in \widetilde{\mathcal{S}}^K$ defined through $\boldsymbol{s}_\kappa = 1$ if $\kappa \in \mathcal{X}$, $= -1$ else, $\kappa \in \mathcal{K}$, attains an objective value of zero in Problem (EC.7). Conversely, if the optimal objective value of Problem (EC.7) is zero, then the set $\mathcal{X} := \{\kappa \in \mathcal{A} : \boldsymbol{s}_\kappa = 1\}$ is such that $\sum_{i\in\mathcal{X}} \boldsymbol{w}_i = W$ and the claim follows. This concludes the proof of the second item.

Both items are thus proved. $\square$

Proof of Proposition 1 Fix $\boldsymbol{\iota} \in \mathcal{C}^K$, $\boldsymbol{s} \in \widetilde{\mathcal{S}}^K$, and $\boldsymbol{x} \in \mathcal{R}$ and consider the inner minimization problem in ($\mathcal{R}_{\text{risk}}$). Since no element of $\boldsymbol{s}$ is zero, this inner problem is expressible as

$$\begin{array}{ll} \text{minimize} & \boldsymbol{u}^\top \boldsymbol{x} \\ \text{subject to} & \boldsymbol{u} \in \mathbb{R}^J \\ & \boldsymbol{u}^\top (\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}) \geq 0 \quad \forall \kappa \in \mathcal{K} : \boldsymbol{s}_\kappa = 1 \\ & \boldsymbol{u}^\top (\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}) \leq 0 \quad \forall \kappa \in \mathcal{K} : \boldsymbol{s}_\kappa = -1 \\ & \boldsymbol{B}\boldsymbol{u} \geq \boldsymbol{b}, \end{array}$$

or equivalently as

$$\begin{array}{ll} \text{minimize} & \boldsymbol{u}^\top \boldsymbol{x} \\ \text{subject to} & \boldsymbol{u} \in \mathbb{R}^J \\ & \boldsymbol{u}^\top [\boldsymbol{s}_\kappa \, (\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa})] \geq 0 \quad \forall \kappa \in \mathcal{K} \\ & \boldsymbol{B}\boldsymbol{u} \geq \boldsymbol{b}. \end{array}$$

Its dual reads

$$\begin{array}{ll} \text{maximize} & \boldsymbol{b}^\top \boldsymbol{\beta} \\ \text{subject to} & \boldsymbol{\alpha} \in \mathbb{R}^K_+, \, \boldsymbol{\beta} \in \mathbb{R}^M_+ \\ & \displaystyle\sum_{\kappa\in\mathcal{K}} \boldsymbol{s}_\kappa \, (\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}) \, \boldsymbol{\alpha}_\kappa + \boldsymbol{B}^\top \boldsymbol{\beta} = \boldsymbol{x}. \end{array} \tag{EC.8}$$

Next, we claim that the primal-dual pair above satisfies strong duality. If $\boldsymbol{s} \in \mathcal{S}(\boldsymbol{\iota})$, then $\mathcal{U}(\boldsymbol{\iota}, \boldsymbol{s})$ and thus also $\widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s})$ are non-empty. Since $\widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s})$ is non-empty and compact, by Assumption 2, it follows that the primal problem above is feasible and bounded, so that it is solvable. We show that the statement also holds if $\boldsymbol{s} \notin \mathcal{S}(\boldsymbol{\iota})$. If $\boldsymbol{s} \notin \mathcal{S}(\boldsymbol{\iota})$, then the primal problem is infeasible. This implies that the dual is either infeasible or unbounded. We show that it cannot be infeasible. Fix $\boldsymbol{\alpha} = \boldsymbol{0}$ in the dual. Then, the dual reduces to

$$\begin{array}{ll} \text{maximize} & \boldsymbol{b}^\top \boldsymbol{\beta} \\ \text{subject to} & \boldsymbol{\beta} \in \mathbb{R}^M_+ \\ & \boldsymbol{B}^\top \boldsymbol{\beta} = \boldsymbol{x}. \end{array}$$

By Farkas' lemma, exactly one of the following alternatives must hold: *a)* There exists $\boldsymbol{\beta} \in \mathbb{R}^M_+$ such that $\boldsymbol{B}^\top \boldsymbol{\beta} = \boldsymbol{x}$; or *b)* There exists $\boldsymbol{u} \in \mathbb{R}^J$ such that $\boldsymbol{B}\boldsymbol{u} \geq \boldsymbol{0}$ and $\boldsymbol{u}^\top \boldsymbol{x} < 0$. Since the set $\mathcal{U}^0$ is bounded, by Assumption 2, its recession cone coincides with the origin implying that assertion *b)* cannot hold so that the dual must be feasible. We conclude that the dual must be unbounded so that the optimal objective values of the primal and dual problems coincide for all $\boldsymbol{s} \in \widetilde{\mathcal{S}}^K$.

Combining the dual problem (EC.8) with the outer maximization in ($\mathcal{R}_{\text{risk}}$) yields Problem (3) which, for fixed $\boldsymbol{\iota} \in \mathcal{C}$, $\boldsymbol{s} \in \widetilde{\mathcal{S}}$, and $\boldsymbol{x} \in \mathcal{R}$ is a linear program if $\mathcal{R}$ is polyhedral and a mixed-binary linear program if $\mathcal{R}$ also involves integrality constraints. $\square$

Proof of Observation 2 Fix $\boldsymbol{\iota} \in \mathcal{C}^K$. Then,

$$\begin{aligned}
&\min_{\boldsymbol{s}\in\widetilde{\mathcal{S}}^K} \ \max_{\boldsymbol{x}\in\mathcal{R}} \ \min_{\boldsymbol{u}\in\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s})} \ \boldsymbol{u}^\top \boldsymbol{x} \\
=\ &\min \left\{ \max_{\boldsymbol{x}^{\boldsymbol{s}}\in\mathcal{R}} \ \min_{\boldsymbol{u}\in\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s})} \ \boldsymbol{u}^\top \boldsymbol{x}^{\boldsymbol{s}} \right\}_{\boldsymbol{s}\in\widetilde{\mathcal{S}}^K} \\
=\ &\max_{\substack{\boldsymbol{x}^{\boldsymbol{s}}\in\mathcal{R}:\\ \boldsymbol{s}\in\widetilde{\mathcal{S}}^K}} \ \min \left\{ \min_{\boldsymbol{u}\in\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s})} \ \boldsymbol{u}^\top \boldsymbol{x}^{\boldsymbol{s}} \right\}_{\boldsymbol{s}\in\widetilde{\mathcal{S}}^K} \\
=\ &\max_{\substack{\boldsymbol{x}^{\boldsymbol{s}}\in\mathcal{R}:\\ \boldsymbol{s}\in\widetilde{\mathcal{S}}^K}} \ \min_{\boldsymbol{s}\in\widetilde{\mathcal{S}}^K} \ \min_{\boldsymbol{u}\in\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s})} \ \boldsymbol{u}^\top \boldsymbol{x}^{\boldsymbol{s}},
\end{aligned} \tag{EC.9}$$

where the second equality follows from the fact that each term in the minimum involves a different set of variables in the maximum, which can be optimized separately. Since the choice of $\boldsymbol{\iota} \in \mathcal{C}^K$ was arbitrary, the claim follows. $\square$

Proof of Lemma 2 Problem (4) can be written in epigraph form equivalently as

$$\begin{array}{ll}
\text{maximize} & \tau \\
\text{subject to} & \boldsymbol{\iota} \in \mathcal{C}^K,\ \boldsymbol{x}^{\boldsymbol{s}} \in \mathcal{R},\ \boldsymbol{s} \in \widetilde{\mathcal{S}}^K \\
& \tau \ \leq \ \displaystyle\min_{\boldsymbol{s}\in\widetilde{\mathcal{S}}^K} \ \min_{\boldsymbol{u}\in\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s})} \ \boldsymbol{u}^\top \boldsymbol{x}^{\boldsymbol{s}}.
\end{array}$$

The problem above is in turn equivalent to

$$\begin{array}{ll}
\text{maximize} & \tau \\
\text{subject to} & \tau \in \mathbb{R},\ \boldsymbol{\iota} \in \mathcal{C}^K,\ \boldsymbol{x}^{\boldsymbol{s}} \in \mathcal{R},\ \boldsymbol{s} \in \widetilde{\mathcal{S}}^K \\
& \tau \ \leq \ \displaystyle\min_{\boldsymbol{u}\in\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s})} \ \boldsymbol{u}^\top \boldsymbol{x}^{\boldsymbol{s}} \quad \forall \boldsymbol{s} \in \widetilde{\mathcal{S}}^K.
\end{array} \tag{EC.10}$$

To reformulate this robust problem which involves infinitely many constraints as a finite program, we employ techniques from robust optimization. Fix $\boldsymbol{s} \in \widetilde{\mathcal{S}}^K$, $\boldsymbol{\iota} \in \mathcal{C}^K$, and $\boldsymbol{x}^{\boldsymbol{s}} \in \mathcal{R}$, and consider the minimization subproblem associated with $\boldsymbol{s}$ in the epigraph constraint. This subproblem reads

$$\begin{array}{rl} \text{minimize} & \boldsymbol{u}^\top \boldsymbol{x}^{\boldsymbol{s}} \\ \text{subject to} & \boldsymbol{u} \in \mathbb{R}^J \\ & \boldsymbol{u}^\top [\boldsymbol{s}_\kappa (\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa})] \geq 0 \quad \forall \kappa \in \mathcal{K} \\ & \boldsymbol{B}\boldsymbol{u} \geq \boldsymbol{b}. \end{array}$$

Its dual is expressible as

$$\begin{array}{rl} \text{maximize} & \boldsymbol{b}^\top \boldsymbol{\beta}^{\boldsymbol{s}} \\ \text{subject to} & \boldsymbol{\alpha}^{\boldsymbol{s}} \in \mathbb{R}_+^K,\ \boldsymbol{\beta}^{\boldsymbol{s}} \in \mathbb{R}_+^M \\ & \displaystyle\sum_{\kappa \in \mathcal{K}} \boldsymbol{s}_\kappa \, (\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa})\, \boldsymbol{\alpha}_\kappa^{\boldsymbol{s}} + \boldsymbol{B}^\top \boldsymbol{\beta}^{\boldsymbol{s}} = \boldsymbol{x}^{\boldsymbol{s}} \end{array}$$

Following a proof strategy similar to that taken in the proof of Proposition 1 (based on Farkas' lemma), we conclude that the optimal objective values of the primal-dual pair above are always equal (even when the primal is infeasible). Replacing each minimization subproblem in Problem (EC.10) with its dual, we obtain the equivalent reformulation

$$\begin{array}{rl} \max & \tau \\ \text{s.t.} & \tau \in \mathbb{R},\ \boldsymbol{\iota} \in \mathcal{C}^K \\ & \left.\begin{array}{l} \boldsymbol{\alpha}^{\boldsymbol{s}} \in \mathbb{R}_+^K,\ \boldsymbol{\beta}^{\boldsymbol{s}} \in \mathbb{R}_+^M,\ \boldsymbol{x}^{\boldsymbol{s}} \in \mathcal{R} \\ \tau \ \leq \ \boldsymbol{b}^\top \boldsymbol{\beta}^{\boldsymbol{s}} \\ \displaystyle\sum_{\kappa \in \mathcal{K}} \boldsymbol{s}_\kappa \, (\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa})\, \boldsymbol{\alpha}_\kappa^{\boldsymbol{s}} + \boldsymbol{B}^\top \boldsymbol{\beta}^{\boldsymbol{s}} = \boldsymbol{x}^{\boldsymbol{s}} \end{array}\right\} \quad \forall \boldsymbol{s} \in \widetilde{\mathcal{S}}^K, \end{array}$$

which concludes the proof. $\square$

Proof of Theorem 2 We begin by showing that Problem (5) is equivalent to

$$
\begin{aligned}
\text{maximize} \quad & \tau \\
\text{subject to} \quad & \tau \in \mathbb{R},\ \boldsymbol{v}^\kappa,\ \boldsymbol{w}^\kappa \in \{0,1\}^I,\ \kappa \in \mathcal{K} \\
& \boldsymbol{\alpha}^{\boldsymbol{s}} \in \mathbb{R}^K_+,\ \boldsymbol{\beta}^{\boldsymbol{s}} \in \mathbb{R}^M_+,\ \boldsymbol{x}^{\boldsymbol{s}} \in \mathcal{R},\ \boldsymbol{s} \in \widetilde{\mathcal{S}}^K \\
& \left.\begin{aligned}
& \tau \ \le\ \boldsymbol{b}^\top \boldsymbol{\beta}^{\boldsymbol{s}} \\
& \sum_{\kappa \in \mathcal{K}} \boldsymbol{s}_\kappa \sum_{i \in \mathcal{I}} \boldsymbol{x}^i \left(\boldsymbol{v}^\kappa_i - \boldsymbol{w}^\kappa_i\right) \boldsymbol{\alpha}^{\boldsymbol{s}}_\kappa + \boldsymbol{B}^\top \boldsymbol{\beta}^{\boldsymbol{s}} \ =\ \boldsymbol{x}^{\boldsymbol{s}}
\end{aligned} \right\} \quad \forall \boldsymbol{s} \in \widetilde{\mathcal{S}}^K \\
& \left.\begin{aligned}
& \mathbf{e}^\top \boldsymbol{v}^\kappa = 1,\ \mathbf{e}^\top \boldsymbol{w}^\kappa = 1 \\
& 1 - \boldsymbol{w}^\kappa_i \ \ge \sum_{i' : i' \ge i} \boldsymbol{v}^\kappa_{i'} \quad \forall i \in \mathcal{I}
\end{aligned} \right\} \quad \forall \kappa \in \mathcal{K}.
\end{aligned}
\tag{EC.11}
$$

For the first direction, let $(\tau, \boldsymbol{\iota}, \{\boldsymbol{\alpha}^{\boldsymbol{s}}, \boldsymbol{\beta}^{\boldsymbol{s}}, \boldsymbol{x}^{\boldsymbol{s}}\}_{\boldsymbol{s} \in \widetilde{\mathcal{S}}^K})$ be feasible in Problem (5). For each $\kappa \in \mathcal{K}$, define $\boldsymbol{v}^\kappa$ and $\boldsymbol{w}^\kappa \in \{0,1\}^I$ through

$$
\boldsymbol{v}^\kappa_i := \begin{cases} 1 & \text{if } \boldsymbol{\iota}^\kappa_1 = i \\ 0 & \text{else,} \end{cases} \qquad \text{and} \qquad \boldsymbol{w}^\kappa_i := \begin{cases} 1 & \text{if } \boldsymbol{\iota}^\kappa_2 = i \\ 0 & \text{else,} \end{cases}
$$

for each $i \in \mathcal{I}$. Fix $\kappa \in \mathcal{K}$. Since $\boldsymbol{\iota}^\kappa \in \mathcal{C}$, it readily follows that $\boldsymbol{\iota}^\kappa_1$ and $\boldsymbol{\iota}^\kappa_2$ are both in $\mathcal{I}$, implying that $\mathbf{e}^\top \boldsymbol{v}^\kappa = \mathbf{e}^\top \boldsymbol{w}^\kappa = 1$. Fix $i \in \mathcal{I}$. If $\boldsymbol{w}^\kappa_i = 0$, then it holds that $1 - \boldsymbol{w}^\kappa_i \ge \sum_{i' : i' \ge i} \boldsymbol{v}^\kappa_{i'}$. On the other hand, if $\boldsymbol{w}^\kappa_i = 1$, then it holds that $\boldsymbol{\iota}^\kappa_2 = i$. But since $\boldsymbol{\iota}^\kappa \in \mathcal{C}$, it holds that $\boldsymbol{\iota}^\kappa_1 < \boldsymbol{\iota}^\kappa_2 = i$, i.e., there must exist $i' < i$ such that $\boldsymbol{\iota}^\kappa_1 = i'$ and $\boldsymbol{v}_{i'} = 1$. This in turn implies that $\sum_{i' : i' \ge i} \boldsymbol{v}^\kappa_{i'} = 0$, and therefore $1 - \boldsymbol{w}^\kappa_i \ge \sum_{i' : i' \ge i} \boldsymbol{v}^\kappa_{i'}$ holds. Since the choice of $\kappa \in \mathcal{K}$ and $i \in \mathcal{I}$ was arbitrary, it holds that

$$
\left.\begin{aligned}
& \mathbf{e}^\top \boldsymbol{v}^\kappa = 1,\ \mathbf{e}^\top \boldsymbol{w}^\kappa = 1 \\
& 1 - \boldsymbol{w}^\kappa_i \ \ge \sum_{i' : i' \ge i} \boldsymbol{v}^\kappa_{i'} \quad \forall i \in \mathcal{I}
\end{aligned} \right\} \quad \forall \kappa \in \mathcal{K}.
$$

Fix $\boldsymbol{s} \in \widetilde{\mathcal{S}}^K$ and $\kappa \in \mathcal{K}$. Then,

$$
\sum_{i \in \mathcal{I}} \boldsymbol{x}^i \left(\boldsymbol{v}^\kappa_i - \boldsymbol{w}^\kappa_i\right) \boldsymbol{\alpha}^{\boldsymbol{s}}_\kappa \ =\ (\boldsymbol{x}^{\boldsymbol{\iota}^\kappa_1} - \boldsymbol{x}^{\boldsymbol{\iota}^\kappa_1}) \boldsymbol{\alpha}^{\boldsymbol{s}}_\kappa
$$

holds trivially by definition of $\boldsymbol{v}^\kappa$ and $\boldsymbol{w}^\kappa$. Thus, $(\tau, \{\boldsymbol{v}^\kappa, \boldsymbol{w}^\kappa\}_{\kappa \in \mathcal{K}}, \{\boldsymbol{\alpha}^{\boldsymbol{s}}, \boldsymbol{\beta}^{\boldsymbol{s}}, \boldsymbol{x}^{\boldsymbol{s}}\}_{\boldsymbol{s} \in \widetilde{\mathcal{S}}^K})$ is feasible in Problem (EC.11) with objective value equal to τ.

For the other direction, let $(\tau, \{\boldsymbol{v}^\kappa, \boldsymbol{w}^\kappa\}_{\kappa \in \mathcal{K}}, \{\boldsymbol{\alpha}^{\boldsymbol{s}}, \boldsymbol{\beta}^{\boldsymbol{s}}, \boldsymbol{x}^{\boldsymbol{s}}\}_{\boldsymbol{s} \in \widetilde{\mathcal{S}}^K})$ be feasible in Problem (EC.11). Define

$$
\boldsymbol{\iota}^\kappa_1 = \sum_{i \in \mathcal{I}} i \cdot \mathbb{I}(\boldsymbol{v}^\kappa_i = 1) \qquad \text{and} \qquad \boldsymbol{\iota}^\kappa_2 = \sum_{i \in \mathcal{I}} i \cdot \mathbb{I}(\boldsymbol{w}^\kappa_i = 1).
$$

Fix $\kappa \in \mathcal{K}$. Then, $\boldsymbol{\iota}^\kappa_1$ and $\boldsymbol{\iota}^\kappa_2$ are both in the set $\mathcal{I}$. It follows from

$$
1 - \boldsymbol{w}^\kappa_i \ \ge \sum_{i' : i' \ge i} \boldsymbol{v}^\kappa_{i'} \quad \forall i \in \mathcal{I}
$$

that if $\boldsymbol{w}_i^\kappa = 1$, then $\boldsymbol{\iota}_2^\kappa = i$ and $\sum_{i':i'\geq i} \boldsymbol{v}_{i'}^\kappa = 0$, implying that $\sum_{i':i'<i} \boldsymbol{v}_{i'}^\kappa = 1$, i.e., $\boldsymbol{\iota}_1^\kappa < i = \boldsymbol{\iota}_2^\kappa$ and it holds that $\boldsymbol{\iota}^\kappa \in \mathcal{C}$. Since the choice of κ was arbitrary, $\boldsymbol{\iota} \in \mathcal{C}^K$. Fix $\boldsymbol{s} \in \widetilde{\mathcal{S}}^K$ and $\kappa \in \mathcal{K}$. Then,

$$(\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa})\, \boldsymbol{\alpha}_\kappa^{\boldsymbol{s}} = \sum_{i\in\mathcal{I}} \boldsymbol{x}^i \left(\boldsymbol{v}_i^\kappa - \boldsymbol{w}_i^\kappa\right) \boldsymbol{\alpha}_\kappa^{\boldsymbol{s}}$$

holds trivially by definition of $\boldsymbol{\iota}^\kappa$. Therefore, $(\tau, \boldsymbol{\iota}, \{\boldsymbol{\alpha}^{\boldsymbol{s}}, \boldsymbol{\beta}^{\boldsymbol{s}}, \boldsymbol{x}^{\boldsymbol{s}}\}_{\boldsymbol{s}\in\widetilde{\mathcal{S}}^K})$ is feasible in Problem (5) with objective value equal to τ. We have thus shown that Problems (5) and (EC.11) are equivalent.

Equivalence of the bilinear problem (EC.11) and the MILP (6) follows directly by using standard linearization techniques which apply since all bilinear terms involve products of binary and real valued variables, see e.g., Hillier (2012). $\square$

Proof of Proposition 2 Since $\boldsymbol{\iota}$ is feasible in Problem $(\mathcal{CCG}_{\text{risk}}^{\text{master}}(\mathcal{S}'))$, it follows that $\boldsymbol{\iota} \in \mathcal{C}^K$. By following a proof strategy similar to that in the Proof of Proposition 1 (based on Farkas' lemma), it follows that $\boldsymbol{\iota}$ is feasible in Problem $(\widetilde{\mathcal{P}}_{\text{off,risk}}^K)$. Thus, it remains to show that for any given $\boldsymbol{\iota} \in \mathcal{C}^K$, Problem $(\mathcal{CCG}_{\text{risk}}^{\text{feas}}(\boldsymbol{\iota}))$ is equivalent to

$$\underset{\boldsymbol{s}\in\widetilde{\mathcal{S}}^K}{\text{minimize}} \quad \max_{\boldsymbol{x}\in\mathcal{R}} \quad \min_{\boldsymbol{u}\in\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s})} \quad \boldsymbol{u}^\top \boldsymbol{x} \tag{EC.12}$$

in the sense that the two problems have the same optimal objective value. To this end, fix $\boldsymbol{\iota} \in \mathcal{C}^K$. Using an epigraph reformulation, we can write Problem (EC.12) equivalently as

$$\begin{array}{ll} \text{minimize} & \theta \\ \text{subject to} & \theta \in \mathbb{R},\ \boldsymbol{s} \in \widetilde{\mathcal{S}}^K \\ & \theta \geq \displaystyle\max_{\boldsymbol{x}\in\mathcal{R}} \ \min_{\boldsymbol{u}\in\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s})} \ \boldsymbol{u}^\top \boldsymbol{x}. \end{array} \tag{EC.13}$$

Since $\mathcal{R}$ has fixed finite cardinality, we can equivalently express the above problem as

$$\begin{array}{ll} \text{minimize} & \theta \\ \text{subject to} & \theta \in \mathbb{R},\ \boldsymbol{s} \in \widetilde{\mathcal{S}}^K \\ & \boldsymbol{u}^{\boldsymbol{x}} \in \widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s}) \quad \forall \boldsymbol{x} \in \mathcal{R} \\ & \theta \geq (\boldsymbol{u}^{\boldsymbol{x}})^\top \boldsymbol{x} \quad \forall \boldsymbol{x} \in \mathcal{R}. \end{array}$$

From the definition of $\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s})$, we can rewrite the preceding problem equivalently as

$$\begin{array}{ll} \text{minimize} & \theta \\ \text{subject to} & \theta \in \mathbb{R},\ \boldsymbol{u}^{\boldsymbol{x}} \in \mathcal{U}^0 \ \ \forall \boldsymbol{x} \in \mathcal{R}, \ \ \boldsymbol{s} \in \overline{\mathcal{S}}^K \\ & \theta \geq (\boldsymbol{u}^{\boldsymbol{x}})^\top \boldsymbol{x} \quad \forall \boldsymbol{x} \in \mathcal{R} \\ & \left.\begin{array}{l} (\boldsymbol{u}^{\boldsymbol{x}})^\top (\boldsymbol{x}^{\boldsymbol{\iota}_k} - \boldsymbol{x}^{\boldsymbol{\iota}'_k}) \geq 0 \quad \forall \kappa \in \mathcal{K} : \boldsymbol{s}_\kappa = 1 \\ (\boldsymbol{u}^{\boldsymbol{x}})^\top (\boldsymbol{x}^{\boldsymbol{\iota}_k} - \boldsymbol{x}^{\boldsymbol{\iota}'_k}) \leq 0 \quad \forall \kappa \in \mathcal{K} : \boldsymbol{s}_\kappa = -1 \end{array}\right\} \quad \forall \boldsymbol{x} \in \mathcal{R}, \end{array}$$

which, for M sufficiently large, is equivalent to Problem ($\mathcal{CCG}^{\text{feas}}_{\text{risk}}(\boldsymbol{\iota})$). Thus, Problems (EC.12) and ($\mathcal{CCG}^{\text{feas}}_{\text{risk}}(\boldsymbol{\iota})$) are equivalent. Since the choice of $\boldsymbol{\iota}$ was arbitrary, the claim follows. $\square$

Proof of Lemma 3 *(i)* By virtue of Proposition 2, θ yields a lower bound to the optimal value of Problem ($\widetilde{\mathcal{P}}^K_{\text{off,risk}}$). At the same time, since $\mathcal{S}' \subseteq \widetilde{\mathcal{S}}^K$, it is evident that τ is an upper bound to the optimal objective value of Problem ($\widetilde{\mathcal{P}}^K_{\text{off,risk}}$). Therefore, $\theta \leq \tau$.

(ii) Suppose that $\theta = \tau$ and that there exists $\boldsymbol{s} \in \widetilde{\mathcal{S}}^K$ such that Problem ($\mathcal{CCG}^{\text{sub},\boldsymbol{s}}_{\text{risk}}(\tau, \boldsymbol{\iota})$) is infeasible. This implies that Problem ($\mathcal{CCG}^{\text{master}}_{\text{risk}}(\mathcal{S}')$) is solvable and that there exists $\boldsymbol{s} \in \widetilde{\mathcal{S}}^K$ such that τ is strictly larger than the optimal objective value of

$$
\begin{array}{ll}
\text{maximize} & \boldsymbol{b}^\top \boldsymbol{\beta} \\
\text{subject to} & \boldsymbol{x} \in \mathcal{R},\ \boldsymbol{\alpha} \in \mathbb{R}^K,\ \boldsymbol{\beta} \in \mathbb{R}^M_+ \\
& \displaystyle\sum_{\kappa \in \mathcal{K}} \boldsymbol{s}_\kappa \ (\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa})\, \boldsymbol{\alpha}_\kappa + \boldsymbol{B}^\top \boldsymbol{\beta} = \boldsymbol{x}.
\end{array}
\tag{EC.14}
$$

An inspection of the proof of Proposition 2 reveals that Problems ($\mathcal{CCG}^{\text{feas}}_{\text{risk}}(\boldsymbol{\iota})$) and (EC.12) are equivalent. An inspection of the proof of Lemma 2 shows that Problem (EC.12) is equivalent to

$$
\begin{array}{lll}
\underset{\boldsymbol{s} \in \widetilde{\mathcal{S}}^K}{\text{minimize}} & \max & \boldsymbol{b}^\top \boldsymbol{\beta} \\
& \text{s.t.} & \boldsymbol{x} \in \mathcal{R},\ \boldsymbol{\alpha} \in \mathbb{R}^K,\ \boldsymbol{\beta} \in \mathbb{R}^M_+ \\
& & \displaystyle\sum_{\kappa \in \mathcal{K}} \boldsymbol{s}_\kappa\, (\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa})\, \boldsymbol{\alpha}_\kappa + \boldsymbol{B}^\top \boldsymbol{\beta} = \boldsymbol{x}.
\end{array}
\tag{EC.15}
$$

Thus, Problems (EC.15) and ($\mathcal{CCG}^{\text{feas}}_{\text{risk}}(\boldsymbol{\iota})$) are equivalent and Problem (EC.15) has an optimal objective value of θ. This implies that $\theta < \tau$, a contradiction.

(iii) Suppose $\theta < \tau$, and let $\boldsymbol{s}$ be defined as in the premise of the lemma. Then, the proof of item *(ii)* reveals that $\boldsymbol{s}$ must be optimal in Problem (EC.15) with associated optimal value θ. This in turn implies that $\boldsymbol{s}$ is such that the optimal objective value of Problem (EC.14) is strictly less than τ, implying that the $\boldsymbol{s}$th subproblem ($\mathcal{CCG}^{\text{sub},\boldsymbol{s}}_{\text{risk}}(\tau, \boldsymbol{\iota})$) is infeasible, which concludes the proof.

We have thus proved all claims. $\square$

Proof of Theorem 3 First, note that finite termination is guaranteed since at each iteration, either UB $-$ LB $\leq \delta$ (in which case the algorithm terminates) or a new set of constraints (indexed by the infeasible index $\boldsymbol{s}$) is added to the master problem $\mathcal{CCG}^{\text{master}}_{\text{risk}}(\mathcal{S}')$, see Lemma 3. Since the set of all indices $\widetilde{\mathcal{S}}^K$ is finite, the algorithm will terminate in a finite number of steps. Second, by construction, at any iteration of the algorithm, τ (i.e., UB) provides a upper bound on the optimal objective value of the problem. On the other hand, the returned (feasible) solution has as objective value θ (i.e., LB). Since the algorithm only terminates if $\mathrm{UB}-\mathrm{LB}\leq\delta$, 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. $\square$

Proof of Theorem 4 The proof of this statement parallels exactly the proof of Theorem 1 part *(b)* and can thus be omitted. $\square$

EC.2. Proofs of Statements in Sections 5

Proof of Lemma 4 The proof parallels exactly the proof of Lemma 1 and Observation 1 and can thus be omitted. $\square$

Proof of Theorem 5 We use a reduction from PARTITION, as in the proof of Theorem 1. We aim to reduce the partition problem to evaluating the objective function of an instance of Problem ($\widetilde{\mathcal{P}}^K_{\text{off,regret}}$). For convenience, we work with the negative of the objective function of Problem ($\widetilde{\mathcal{P}}^K_{\text{off,regret}}$) given by

$$\underset{\boldsymbol{s}\in\widetilde{\mathcal{S}}^K}{\text{minimize}}\quad \max_{\boldsymbol{x}\in\mathcal{R}}\quad \min_{\boldsymbol{u}\in\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s})}\quad \left\{\min_{\boldsymbol{x}'\in\mathcal{R}}\ \boldsymbol{u}^\top\boldsymbol{x}-\boldsymbol{u}^\top\boldsymbol{x}'\right\}.$$

To this end, fix an instance $(n,\boldsymbol{w},W)$ of the partition problem. Set $J:=2n+3$, $K:=n$, and

$$\mathcal{U}^0 := \left\{\begin{array}{rl}\boldsymbol{u}\in\mathbb{R}^{2n+3}: & \boldsymbol{u}_i\in[0,1]\quad\forall i\in\{1,\ldots,n\}\\ & \boldsymbol{u}_{n+i} = \boldsymbol{u}_i-0.5\quad\forall i\in\{1,\ldots,n\}\\ & \boldsymbol{u}_{2n+1}\in[0,1]\\ & \boldsymbol{u}_{2n+2}\in[-1,1]\\ & \boldsymbol{u}_{2n+3}=0.5\end{array}\right\}.$$

Also, for each $\kappa\in\mathcal{K}$, let $\boldsymbol{\iota}_1^\kappa:=\mathbf{e}_\kappa$ and $\boldsymbol{\iota}_2^\kappa:=\mathbf{e}_{2n+3}$. Then, given a choice $\boldsymbol{s}\in\widetilde{\mathcal{S}}^K$, we have

$$\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s}):=\left\{\begin{array}{rl}\boldsymbol{u}\in\mathcal{U}^0: & \boldsymbol{u}_\kappa\geq 0.5\quad\forall\kappa\in\mathcal{K}:\boldsymbol{s}_\kappa=1\\ & \boldsymbol{u}_\kappa\leq 0.5\quad\forall\kappa\in\mathcal{K}:\boldsymbol{s}_\kappa=-1\end{array}\right\}.$$

Finally, we define the recommendation set $\mathcal{R}\subseteq\mathbb{R}^J$ through $\mathcal{R}:=\{\boldsymbol{x}^1,\boldsymbol{x}^2\}$, where $\boldsymbol{x}^1:=(\boldsymbol{w}^\top,\mathbf{0}^\top,-W,0,0)^\top$ and $\boldsymbol{x}^2:=(\mathbf{0}^\top,-\boldsymbol{w}^\top,W,2W,0)^\top$. For any fixed $\boldsymbol{s}\in\widetilde{\mathcal{S}}^K$, define

$$Z(\boldsymbol{s})\quad:=\quad\max_{\boldsymbol{x}\in\mathcal{R}}\quad\min_{\boldsymbol{u}\in\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s})}\quad\min_{\boldsymbol{x}'\in\mathcal{R}}\ \left\{\boldsymbol{u}^\top\boldsymbol{x}-\boldsymbol{u}^\top\boldsymbol{x}'\right\}.$$

Then, we have

$$\begin{aligned}
Z(\boldsymbol{s}) &= \max_{\boldsymbol{x}\in\mathcal{R}} \min\left\{ \min_{\boldsymbol{u}\in\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s})} \boldsymbol{u}^\top(\boldsymbol{x}-\boldsymbol{x}^1),\ \min_{\boldsymbol{u}\in\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s})} \boldsymbol{u}^\top(\boldsymbol{x}-\boldsymbol{x}^2)\right\} \\
&= \max\left[\min\left\{ \min_{\boldsymbol{u}\in\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s})} \boldsymbol{u}^\top(\boldsymbol{x}^1-\boldsymbol{x}^1),\ \min_{\boldsymbol{u}\in\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s})} \boldsymbol{u}^\top(\boldsymbol{x}^1-\boldsymbol{x}^2)\right\},\right. \\
&\qquad\qquad \left.\min\left\{ \min_{\boldsymbol{u}\in\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s})} \boldsymbol{u}^\top(\boldsymbol{x}^2-\boldsymbol{x}^1),\ \min_{\boldsymbol{u}\in\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s})} \boldsymbol{u}^\top(\boldsymbol{x}^2-\boldsymbol{x}^2)\right\}\right] \\
&= \max\left[\min\left\{0,\ \min_{\boldsymbol{u}\in\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s})} \boldsymbol{u}^\top(\boldsymbol{x}^1-\boldsymbol{x}^2)\right\},\ \min\left\{ \min_{\boldsymbol{u}\in\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s})} \boldsymbol{u}^\top(\boldsymbol{x}^2-\boldsymbol{x}^1),\ 0\right\}\right].
\end{aligned}$$

Next, observe that

$$\begin{aligned}
&\min_{\boldsymbol{u}\in\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s})} \boldsymbol{u}^\top(\boldsymbol{x}^1-\boldsymbol{x}^2) \\
&= \min_{\boldsymbol{u}\in\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s})} \sum_{i=1}^n \boldsymbol{u}_i\boldsymbol{w}_i + \sum_{i=1}^n \boldsymbol{u}_{n+i}\boldsymbol{w}_i - 2\boldsymbol{u}_{2n+1}W - 2\boldsymbol{u}_{2n+2}W + 0\cdot\boldsymbol{u}_{2n+3} \\
&= \min_{\boldsymbol{u}\in\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s})} \sum_{i=1}^n (2\boldsymbol{u}_i - 0.5)\boldsymbol{w}_i - 2W - 2W \\
&= \min_{\boldsymbol{u}\in\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s})} \sum_{i=1}^n 2\boldsymbol{u}_i\boldsymbol{w}_i - W - 2W - 2W \\
&= \sum_{\kappa\in\mathcal{K}:\boldsymbol{s}_\kappa=1} \boldsymbol{w}_\kappa - W - 4W \\
&\le 0.
\end{aligned}$$

Similarly, it holds that

$$\begin{aligned}
&\min_{\boldsymbol{u}\in\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s})} \boldsymbol{u}^\top(\boldsymbol{x}^2-\boldsymbol{x}^1) \\
&= \min_{\boldsymbol{u}\in\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s})} \sum_{i=1}^n -\boldsymbol{u}_i\boldsymbol{w}_i + \sum_{i=1}^n -\boldsymbol{u}_{n+i}\boldsymbol{w}_i + 2\boldsymbol{u}_{2n+1}W + 2\boldsymbol{u}_{2n+2}W - 0\cdot\boldsymbol{u}_{2n+3} \\
&= \min_{\boldsymbol{u}\in\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s})} \sum_{i=1}^n -(2\boldsymbol{u}_i - 0.5)\boldsymbol{w}_i + 0 - 2W \\
&= \min_{\boldsymbol{u}\in\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s})} \sum_{i=1}^n -2(\boldsymbol{u}_i - 0.5)\boldsymbol{w}_i - W + 0 - 2W - 2W + 2W \\
&= \sum_{\kappa\in\mathcal{K}:\boldsymbol{s}_\kappa=1} -\boldsymbol{w}_\kappa + W - 4W \\
&\le 0.
\end{aligned}$$

Therefore,

$$\begin{aligned}
Z(\boldsymbol{s}) &= \max\left\{ \sum_{\kappa\in\mathcal{K}:\boldsymbol{s}_\kappa=1} \boldsymbol{w}_\kappa - W,\ \sum_{\kappa\in\mathcal{K}:\boldsymbol{s}_\kappa=1} -\boldsymbol{w}_\kappa + W\right\} - 4W \\
&\ge -4W.
\end{aligned}$$

Now, we claim that we are given a "yes" instance of Partition if and only the objective value of $\boldsymbol{\iota}$ in the constructed instance of Problem ($\widetilde{\mathcal{P}}^K_{\text{off,risk}}$) is $-4W$. To this end, note that the objective value of $\boldsymbol{\iota}$ is given by

$$\underset{\boldsymbol{s}\in\widetilde{\mathcal{S}}^K}{\text{minimize}} \quad Z(\boldsymbol{s}) \tag{EC.16}$$

and is lower bounded by $-4W$. If there exists $\mathcal{X} \subset \mathcal{A}$ such that $\sum_{i\in\mathcal{X}} \boldsymbol{w}_i = W$, then the solution $\boldsymbol{s} \in \widetilde{\mathcal{S}}^K$ defined through $\boldsymbol{s}_\kappa = 1$ if $\kappa \in \mathcal{X}$, $= -1$ else, $\kappa \in \mathcal{K}$, attains an objective value of $-4W$ in Problem (EC.16). Conversely, if the optimal objective value of Problem (EC.16) is $-4W$, then the set $\mathcal{X} := \{\kappa \in \mathcal{A} : \boldsymbol{s}_\kappa = 1\}$ is such that $\sum_{i\in\mathcal{X}} \boldsymbol{w}_i = W$ and the claim follows. This concludes the proof of the second item. $\square$

Proof of Proposition 3 Fix $\boldsymbol{\iota} \in \mathcal{C}^K$, $\boldsymbol{s} \in \widetilde{\mathcal{S}}^K$ and consider the resulting regret averse recommendation problem problem given by

$$\underset{\boldsymbol{x}\in\mathcal{R}}{\text{minimize}} \quad \max_{\boldsymbol{u}\in\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s})} \quad \left\{ \max_{\boldsymbol{x}'\in\mathcal{R}} \ \boldsymbol{u}^\top \boldsymbol{x}' - \boldsymbol{u}^\top \boldsymbol{x} \right\}.$$

This problem can be written equivalently in epigraph form as

$$\begin{array}{ll} \text{minimize} & \theta \\ \text{subject to} & \theta \in \mathbb{R},\ \boldsymbol{x} \in \mathcal{R} \\ & \theta \geq \displaystyle\max_{\boldsymbol{u}\in\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s})} \left\{ \max_{\boldsymbol{x}'\in\mathcal{R}} \ \boldsymbol{u}^\top \boldsymbol{x}' - \boldsymbol{u}^\top \boldsymbol{x} \right\}, \end{array}$$

which is in turn equivalent to

$$\begin{array}{ll} \text{minimize} & \theta \\ \text{subject to} & \theta \in \mathbb{R},\ \boldsymbol{x} \in \mathcal{R} \\ & \theta \geq \displaystyle\max_{\boldsymbol{u}\in\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s})} \ \boldsymbol{u}^\top \boldsymbol{x}' - \boldsymbol{u}^\top \boldsymbol{x} \qquad \forall \boldsymbol{x}' \in \mathcal{R}. \end{array} \tag{EC.17}$$

For any fixed $\boldsymbol{x}' \in \mathcal{R}$, the subproblem in the constraints of the above problem can be written explicitly as

$$\begin{array}{ll} \text{maximize} & \boldsymbol{u}^\top (\boldsymbol{x}' - \boldsymbol{x}) \\ \text{subject to} & \boldsymbol{u} \in \mathbb{R}^J \\ & \boldsymbol{u}^\top [\boldsymbol{s}_\kappa \, (\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa})] \geq 0 \qquad \forall \kappa \in \mathcal{K} \\ & \boldsymbol{B}\boldsymbol{u} \geq \boldsymbol{b}. \end{array}$$

Its dual reads

$$\begin{array}{ll} \text{minimize} & \boldsymbol{b}^\top \boldsymbol{\beta} \\ \text{subject to} & \boldsymbol{\alpha} \in \mathbb{R}^K_-,\ \boldsymbol{\beta} \in \mathbb{R}^M_- \\ & \displaystyle\sum_{\kappa\in\mathcal{K}} \boldsymbol{s}_\kappa \, (\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}) \, \boldsymbol{\alpha}_\kappa + \boldsymbol{B}^\top \boldsymbol{\beta} = \boldsymbol{x}' - \boldsymbol{x}. \end{array}$$

By following a reasoning similar to that taken in the Proof of Proposition 1, it can be shown that the primal-dual pair above satisfies strong duality. Combining the dual problem above with the outer minimization in (EC.17) yields the equivalent formulation

$$
\begin{array}{ll}
\text{minimize} & \theta \\
\text{subject to} & \theta \in \mathbb{R},\ \boldsymbol{x} \in \mathcal{R} \\
& \left.\begin{array}{l}
\boldsymbol{\alpha}^{\boldsymbol{x}'} \in \mathbb{R}^K_-,\ \boldsymbol{\beta}^{\boldsymbol{x}'} \in \mathbb{R}^M_- \\
\theta \geq \boldsymbol{b}^\top \boldsymbol{\beta}^{\boldsymbol{x}'} \\
\displaystyle\sum_{\kappa \in \mathcal{K}} \boldsymbol{s}_\kappa \left(\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}\right) \boldsymbol{\alpha}^{\boldsymbol{x}'}_\kappa + \boldsymbol{B}^\top \boldsymbol{\beta}^{\boldsymbol{x}'} = \boldsymbol{x}' - \boldsymbol{x}
\end{array}\right\} \quad \forall \boldsymbol{x}' \in \mathcal{R},
\end{array}
$$

which, for fixed $\boldsymbol{\iota} \in \mathcal{C}$ and $\boldsymbol{s} \in \widetilde{\mathcal{S}}$ is a mixed-binary linear program if $\mathcal{R}$ has fixed finite cardinality. $\square$

Proof of Lemma 5 In a way that parallels the proof of Observation 2, it can be readily shown that Problem (7) is equivalent to

$$
\underset{\boldsymbol{\iota} \in \mathcal{C}^K}{\text{minimize}} \quad \min_{\substack{\boldsymbol{x}^{\boldsymbol{s}} \in \mathcal{R}: \\ \boldsymbol{s} \in \widetilde{\mathcal{S}}^K}} \quad \max_{\boldsymbol{s} \in \widetilde{\mathcal{S}}^K} \quad \max_{\boldsymbol{u} \in \widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s})} \left\{ \max_{\boldsymbol{x}' \in \mathcal{R}} \boldsymbol{u}^\top \boldsymbol{x}' - \boldsymbol{u}^\top \boldsymbol{x}^{\boldsymbol{s}} \right\}.
$$

The above problem can be written in epigraph form equivalently as

$$
\begin{array}{ll}
\text{minimize} & \tau \\
\text{subject to} & \tau \in \mathbb{R},\ \boldsymbol{\iota} \in \mathcal{C}^K,\ \boldsymbol{x}^{\boldsymbol{s}} \in \mathcal{R},\ \boldsymbol{s} \in \widetilde{\mathcal{S}}^K \\
& \tau \ \geq \displaystyle\max_{\boldsymbol{s} \in \widetilde{\mathcal{S}}^K} \ \max_{\boldsymbol{u} \in \widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s})} \ \max_{\boldsymbol{x}' \in \mathcal{R}} \boldsymbol{u}^\top (\boldsymbol{x}' - \boldsymbol{x}^{\boldsymbol{s}}),
\end{array}
$$

which is in turn equivalent to

$$
\begin{array}{ll}
\text{minimize} & \tau \\
\text{subject to} & \tau \in \mathbb{R},\ \boldsymbol{\iota} \in \mathcal{C}^K,\ \boldsymbol{x}^{\boldsymbol{s}} \in \mathcal{R},\ \boldsymbol{s} \in \widetilde{\mathcal{S}}^K \\
& \tau \ \geq \displaystyle\max_{\boldsymbol{u} \in \widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s})} \boldsymbol{u}^\top (\boldsymbol{x}' - \boldsymbol{x}^{\boldsymbol{s}}) \quad \forall \boldsymbol{s} \in \widetilde{\mathcal{S}}^K,\ \forall \boldsymbol{x}' \in \mathcal{R}.
\end{array} \tag{EC.18}
$$

Fix $\boldsymbol{s} \in \widetilde{\mathcal{S}}^K$, $\tau \in \mathbb{R}$, $\boldsymbol{\iota} \in \mathcal{C}^K$, $\boldsymbol{x}^{\boldsymbol{s}} \in \mathcal{R}$, and $\boldsymbol{x}' \in \mathcal{R}$ and consider the associated maximization subproblem in the constraints of the above problem. This reads

$$
\begin{array}{ll}
\text{maximize} & \boldsymbol{u}^\top (\boldsymbol{x}' - \boldsymbol{x}^{\boldsymbol{s}}) \\
\text{subject to} & \boldsymbol{u} \in \mathbb{R}^J \\
& \boldsymbol{u}^\top (\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}) \geq 0 \quad \forall \kappa \in \mathcal{K} : \boldsymbol{s}_\kappa = 1 \\
& \boldsymbol{u}^\top (\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}) \leq 0 \quad \forall \kappa \in \mathcal{K} : \boldsymbol{s}_\kappa = -1 \\
& \boldsymbol{B}\boldsymbol{u} \geq \boldsymbol{b}.
\end{array}
$$

Its dual is expressible as

$$\begin{array}{ll} \text{minimize} & \boldsymbol{b}^\top \boldsymbol{\beta}^{(\boldsymbol{x}',\boldsymbol{s})} \\ \text{subject to} & \boldsymbol{\alpha}^{(\boldsymbol{x}',\boldsymbol{s})} \in \mathbb{R}^K_-,\ \boldsymbol{\beta}^{(\boldsymbol{x}',\boldsymbol{s})} \in \mathbb{R}^M_- \\ & \displaystyle\sum_{\kappa\in\mathcal{K}} \boldsymbol{s}_\kappa \left(\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}\right) \boldsymbol{\alpha}_\kappa^{(\boldsymbol{x}',\boldsymbol{s})} + \boldsymbol{B}^\top \boldsymbol{\beta}^{\boldsymbol{s}} = \boldsymbol{x}' - \boldsymbol{x}^{\boldsymbol{s}}. \end{array}$$

Following a proof strategy similar to that taken in the proof of Proposition 1, it can be shown that the optimal objective values of the primal-dual pair above are always equal (even when the primal is infeasible). Replacing each maximization subproblem in Problem (EC.18) with its dual yields the equivalent formulation

$$\begin{array}{ll} \text{minimize} & \tau \\ \text{subject to} & \tau \in \mathbb{R},\ \boldsymbol{\iota} \in \mathcal{C}^K,\ \boldsymbol{x}^{\boldsymbol{s}} \in \mathcal{R},\ \boldsymbol{s} \in \widetilde{\mathcal{S}}^K \\ & \left.\begin{array}{l} \boldsymbol{\alpha}^{(\boldsymbol{x}',\boldsymbol{s})} \in \mathbb{R}^K_-,\ \boldsymbol{\beta}^{(\boldsymbol{x}',\boldsymbol{s})} \in \mathbb{R}^M_- \\ \tau \ \geq \ \boldsymbol{b}^\top \boldsymbol{\beta}^{(\boldsymbol{x}',\boldsymbol{s})} \\ \displaystyle\sum_{\kappa\in\mathcal{K}} \boldsymbol{s}_\kappa \left(\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}\right) \boldsymbol{\alpha}_\kappa^{(\boldsymbol{x}',\boldsymbol{s})} + \boldsymbol{B}^\top \boldsymbol{\beta}^{(\boldsymbol{x}',\boldsymbol{s})} = \boldsymbol{x}' - \boldsymbol{x}^{\boldsymbol{s}} \end{array}\right\} \quad \forall \boldsymbol{s} \in \widetilde{\mathcal{S}}^K,\ \forall \boldsymbol{x}' \in \mathcal{R}. \end{array} \tag{EC.19}$$

This concludes the proof. $\square$

Proof of Theorem 6 The proof follows directly by following an approach similar to that taken in the Proof of Theorem 2 and is thus omitted. $\square$

Proof of Proposition 4 Since $\boldsymbol{\iota}$ is feasible in Problem ($\mathcal{CCG}^{\text{master}}_{\text{regret}}(\mathcal{S}')$), it follows that $\boldsymbol{\iota} \in \mathcal{C}^K$. Thus, $\boldsymbol{\iota}$ is feasible in Problem ($\widetilde{\mathcal{P}}^K_{\text{off,regret}}$). It remains to show that for any given $\boldsymbol{\iota} \in \mathcal{C}^K$, Problems ($\mathcal{CCG}^{\text{feas}}_{\text{regret}}(\boldsymbol{\iota})$) and

$$\max_{\boldsymbol{s}\in\widetilde{\mathcal{S}}^K} \ \min_{\boldsymbol{x}\in\mathcal{R}} \ \max_{\boldsymbol{u}\in\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s})} \ \left\{ \max_{\boldsymbol{x}'\in\mathcal{R}} \ \boldsymbol{u}^\top \boldsymbol{x}' - \boldsymbol{u}^\top \boldsymbol{x} \right\} \tag{EC.20}$$

have the same optimal objective value. To this end, fix $\boldsymbol{\iota} \in \mathcal{C}^K$. Using an epigraph reformulation, we can write Problem (EC.20) equivalently as

$$\begin{array}{ll} \text{maximize} & \theta \\ \text{subject to} & \theta \in \mathbb{R},\ \boldsymbol{s} \in \widetilde{\mathcal{S}}^K \\ & \theta \ \leq \ \displaystyle\min_{\boldsymbol{x}\in\mathcal{R}} \ \max_{\boldsymbol{u}\in\widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s})} \ \max_{\boldsymbol{x}'\in\mathcal{R}} \ \boldsymbol{u}^\top \boldsymbol{x}' - \boldsymbol{u}^\top \boldsymbol{x}, \end{array} \tag{EC.21}$$

and thus as

$$\begin{array}{ll} \text{maximize} & \theta \\ \text{subject to} & \theta \in \mathbb{R},\ \boldsymbol{s} \in \widetilde{\mathcal{S}}^K, \\ & \boldsymbol{x}'^{,\boldsymbol{x}} \in \mathcal{R},\ \boldsymbol{u}^{\boldsymbol{x}} \in \widetilde{\mathcal{U}}(\boldsymbol{\iota},\boldsymbol{s}) \quad \forall \boldsymbol{x} \in \mathcal{R} \\ & \theta \ \leq \ (\boldsymbol{u}^{\boldsymbol{x}})^\top (\boldsymbol{x}'^{,\boldsymbol{x}} - \boldsymbol{x}) \qquad \forall \boldsymbol{x} \in \mathcal{R}. \end{array}$$

By definition of $\widetilde{\mathcal{U}}(\boldsymbol{\iota}, \boldsymbol{s})$, the above problem can be written as

$$
\begin{array}{lll}
\text{maximize} & \theta & \\
\text{subject to} & \theta \in \mathbb{R},\ \boldsymbol{s} \in \widetilde{\mathcal{S}}^K & \\
& \boldsymbol{x}^{\prime,\boldsymbol{x}} \in \mathcal{R},\ \boldsymbol{u}^{\boldsymbol{x}} \in \mathbb{R}^J \quad \forall \boldsymbol{x} \in \mathcal{R} & \\
& \theta \leq (\boldsymbol{u}^{\boldsymbol{x}})^\top (\boldsymbol{x}^{\prime,\boldsymbol{x}} - \boldsymbol{x}) \quad \forall \boldsymbol{x} \in \mathcal{R} & \\
& \left.\begin{array}{l} (\boldsymbol{u}^{\boldsymbol{x}})^\top (\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}) \geq 0 \quad \forall \kappa \in \mathcal{K} : \boldsymbol{s}_\kappa = 1 \\ (\boldsymbol{u}^{\boldsymbol{x}})^\top (\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}) \leq 0 \quad \forall \kappa \in \mathcal{K} : \boldsymbol{s}_\kappa = -1 \end{array}\right\} & \forall \boldsymbol{x} \in \mathcal{R}.
\end{array}
$$

The claim then follows by rewriting the logical constraints above as linear constraints using a "big-M" constant. $\square$

Proof of Lemma 6 *(i)* By virtue of Proposition 4, θ yields an upper bound to the optimal value of the Problem ($\widetilde{\mathcal{P}}^K_{\text{off,regret}}$). At the same time, since $\widetilde{\mathcal{R}} \subseteq \mathcal{R}$ and $\mathcal{S}' \subseteq \widetilde{\mathcal{S}}^K$, it is evident that τ is a lower bound on the optimal objective value of Problem ($\widetilde{\mathcal{P}}^K_{\text{off,regret}}$). Therefore, $\theta \geq \tau$.

(ii) Suppose that $\theta = \tau$ and that there exists $\boldsymbol{x}' \in \mathcal{R}$ and $\boldsymbol{s} \in \widetilde{\mathcal{S}}^K$ such that Problem ($\mathcal{CCG}^{\text{sub},(\boldsymbol{x}',\boldsymbol{s})}_{\text{regret}}(\tau, \boldsymbol{\iota})$) is infeasible. This implies that Problem ($\mathcal{CCG}^{\text{master}}_{\text{risk}}(\mathcal{S}')$) is solvable and that τ is strictly smaller than the optimal objective value of

$$
\begin{array}{ll}
\text{minimize} & \boldsymbol{b}^\top \boldsymbol{\beta} \\
\text{subject to} & \boldsymbol{\alpha} \in \mathbb{R}^K_-,\ \boldsymbol{\beta} \in \mathbb{R}^M_-,\ \boldsymbol{x} \in \mathcal{R} \\
& \displaystyle\sum_{\kappa \in \mathcal{K}} \boldsymbol{s}_\kappa \, (\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}) \, \boldsymbol{\alpha}_\kappa + \boldsymbol{B}^\top \boldsymbol{\beta} \, = \, \boldsymbol{x}' - \boldsymbol{x}.
\end{array} \tag{EC.22}
$$

An inspection of the proof of Proposition 4 reveals that Problems ($\mathcal{CCG}^{\text{feas}}_{\text{regret}}(\boldsymbol{\iota})$) and (EC.20) are equivalent. An inspection of the proof of Lemma 5 shows that Problem (EC.20) is equivalent to

$$
\begin{array}{lll}
\underset{\boldsymbol{x}' \in \mathcal{R},\ \boldsymbol{s} \in \widetilde{\mathcal{S}}^K}{\text{maximize}} & \min & \boldsymbol{b}^\top \boldsymbol{\beta} \\
& \text{s.t.} & \boldsymbol{\alpha} \in \mathbb{R}^K_-,\ \boldsymbol{\beta} \in \mathbb{R}^M_-,\ \boldsymbol{x} \in \mathcal{R} \\
& & \displaystyle\sum_{\kappa \in \mathcal{K}} \boldsymbol{s}_\kappa \, (\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}) \, \boldsymbol{\alpha}_\kappa + \boldsymbol{B}^\top \boldsymbol{\beta} = \boldsymbol{x}' - \boldsymbol{x}.
\end{array} \tag{EC.23}
$$

Thus, Problems ($\mathcal{CCG}^{\text{feas}}_{\text{regret}}(\boldsymbol{\iota})$) and (EC.23) are equivalent and Problem (EC.23) has an optimal value of θ. This implies that $\tau < \theta$, a contradiction.

(iii) Suppose $\theta > \tau$ and let $(\boldsymbol{x}', \boldsymbol{s})$ be defined as in the premise of the lemma. Then, the proof of item *(ii)* reveals that $\boldsymbol{s}$ is optimal in (EC.20) with associated optimal value θ. This implies that $\boldsymbol{s}$ is such

that the optimal objective value of Problem (EC.22) is strictly greater than τ, so that the $(\boldsymbol{x}', \boldsymbol{s})$th subproblem $(\mathcal{CCG}_{\text{regret}}^{\text{sub},(\boldsymbol{x}',\boldsymbol{s})}(\tau, \boldsymbol{\iota}))$ is infeasible, which concludes the proof.

All claims are thus proved. $\square$

Proof of Theorem 7 The proof mirrors the proof of Theorem 3 and is thus omitted in the interest of space. $\square$

EC.3. Proofs of Statements in Section 6

Proof of Proposition 5 The proof parallels exactly the proof of Proposition 1 and is thus omitted.

Proof of Lemma 7 The proof parallels exactly the proof of Lemma 2 and is thus omitted.

Proof of Theorem 9 The proof parallels the proofs of Proposition 2, Lemma 3, and Theorem 3 and can be omitted.

Proof of Lemma 8 The proof parallels exactly the proof of Lemma 5 and is thus omitted.

Proof of Proposition 6 The proof parallels exactly the proof of Proposition 3 and is thus omitted.

Proof of Theorem 11 The proof parallels the proofs of Proposition 4, Lemma 6, and Theorem 7 and can be omitted.

EC.4. Companion to Section 7

Algorithm 5 details the procedure for generating warm-starts proposed in Section 7.2 to ensure that the lexicographic constraints from Section 7.1 are satisfied.

EC.5. Companion to Section 8

EC.5.1. Additional Numerical Results for Online Elicitation in Section 8

In this section we provide additional results associated with the online preference elicitation numericals on synthetic data that we presented in Section 8.4. These additional results are provided in Figure EC.1. In particular, the two figures show the worst-case true rank and worst-case true utility of the recommended item for the risk averse and regret averse cases, respectively. From the figures, it can be seen that our approaches consistently outperform the state of the art in terms of both worst-case rank and worst-case utility of the recommended item (although these are not quantities that we explicitly optimize for).

Algorithm 5: An algorithm for building warm-starts

Inputs: A feasible solution $(\tau, \boldsymbol{v}, \boldsymbol{w}, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{x}, \overline{\boldsymbol{v}}, \overline{\boldsymbol{w}})$ to Problem (6) with K queries with objective τ;

Output: A warm start $(\tilde{\boldsymbol{v}}, \tilde{\boldsymbol{w}}, \tilde{\boldsymbol{x}})$ to Problem (6) with $K+1$ queries with objective $\geq \tau$;

Initialization:

for $\kappa \in \mathcal{K}$ **do**

  $\boldsymbol{\iota}_1^\kappa \leftarrow \sum_{i \in \mathcal{I}} i \cdot \mathbb{I}(\boldsymbol{v}_i^\kappa = 1)$;

  $\boldsymbol{\iota}_2^\kappa \leftarrow \sum_{i \in \mathcal{I}} i \cdot \mathbb{I}(\boldsymbol{w}_i^\kappa = 1)$;

end

$\widetilde{\mathcal{C}} \leftarrow \{\, \boldsymbol{\iota} \in \mathcal{C} \,:\, \boldsymbol{\iota} \neq \boldsymbol{\iota}^\kappa,\ \kappa \in \mathcal{K}\,\}$;

Select $\boldsymbol{\iota}^{K+1}$ at random from $\widetilde{\mathcal{C}}$;

for $\boldsymbol{s} \in \mathcal{S}^{K+1}$ **do**

  $\tilde{\boldsymbol{x}}^{\boldsymbol{s}} := \boldsymbol{x}^{(\boldsymbol{s}_1, \ldots, \boldsymbol{s}_K)}$;

end

$(\tilde{\boldsymbol{\iota}}, \boldsymbol{\ell}) \leftarrow$ lexicographic order$(\boldsymbol{\iota})$;

for $\boldsymbol{s} \in \mathcal{S}^{K+1}$ **do**

  $\tilde{\boldsymbol{x}}^{\boldsymbol{s}} \leftarrow \tilde{\boldsymbol{x}}^{\boldsymbol{s}(\boldsymbol{\ell})}$;

end

for $\kappa \in \mathcal{K},\ i \in \mathcal{I}$ **do**

  $\tilde{\boldsymbol{v}}_i^\kappa \leftarrow \mathbb{I}(\boldsymbol{\iota}_1^\kappa = i)$;

  $\tilde{\boldsymbol{w}}_i^\kappa \leftarrow \mathbb{I}(\boldsymbol{\iota}_2^\kappa = i)$;

end

Result: $(\tilde{\boldsymbol{v}}, \tilde{\boldsymbol{w}}, \tilde{\boldsymbol{x}})$

Note: The lexicographic order function takes as input a collection of vectors and returns two elements; the first element corresponds to the lexicographically ordered collection of vectors and the second element is a vector of the same dimension as the number of vectors input whose ith element denotes the position of the ith output vector in the input;

EC.5.2. Greedy Solution Approach to Speed-Up Computations in Section 8

To speed-up solution of our problems so as to be able to showcase performance on a variety of instances, we employ a heuristic approach in our experiments in Section 8, as detailed in Algorithm 6. A variant of this approach has been previously used by Vayanos et al. (2019), see also Subramanyam et al. (2017). This algorithm returns a feasible but potentially suboptimal solution to the MBLP counterpart of the offline preference elicitation problem to be solved.

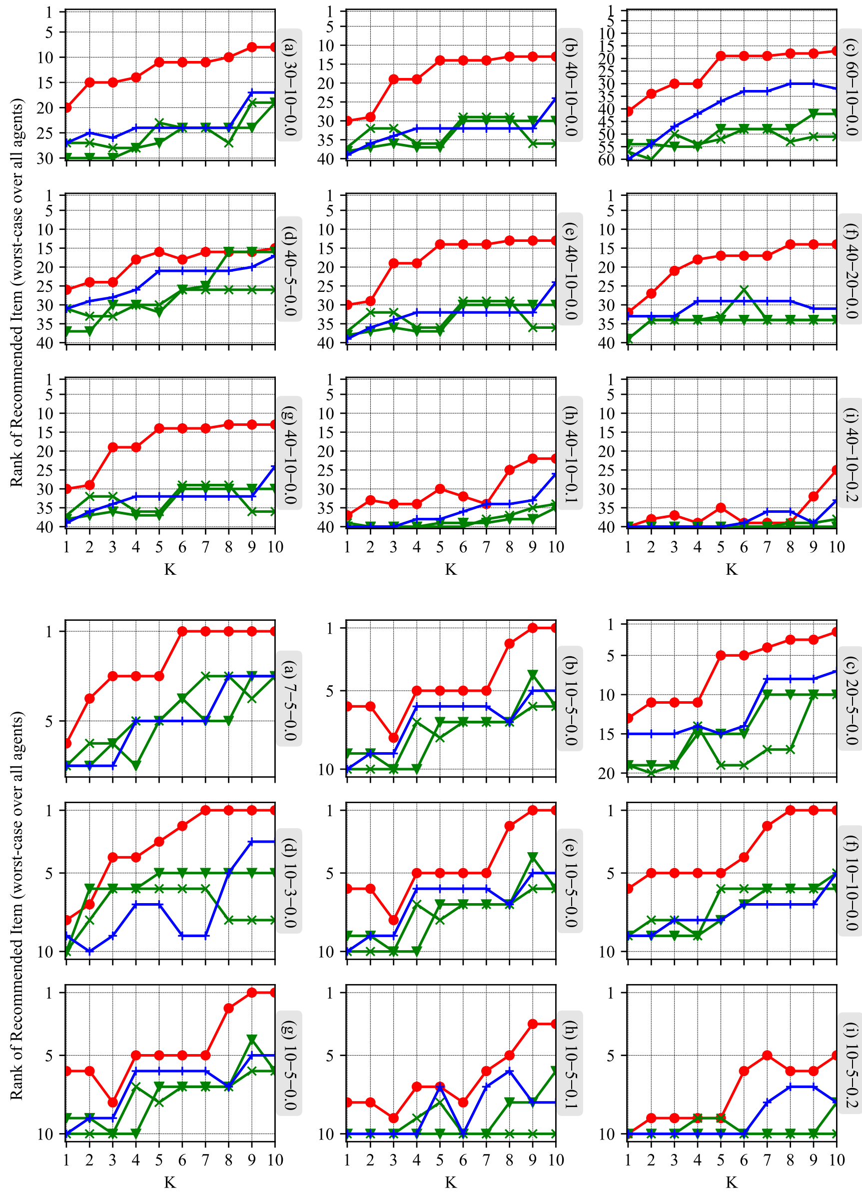

Figure EC.1 The figure at the top (bottom) shows the normalized worst-case true rank of the recommended item (worst-case taken across agents) for the online risk (regret) averse active preference elicitation problem ($\mathcal{P}^K_{\text{on,risk}}$) (($\mathcal{P}^K_{\text{on,regret}}$)) on synthetic data. The facet labels have the same interpretation as in Figure 4 (6). The performance of `E-ON-MMU/R-MMU` (`E-ON-MMR/R-MMR`) across 50 random utility vectors $\boldsymbol{u}^\star$ is shown with red dots. The performance of `E-RAND/R-MMU` (`E-RAND/R-MMR`) is shown with blue crosses. The performance of `E-AC/R-AC` is shown with green crosses. The performance of `E-AC/R-MMU` (`E-AC/R-MMR`) is shown with green triangles.

Algorithm 6: Heuristic algorithm for solving Problem $(\mathcal{P}^{K}_{\text{off,risk}})$, $(\mathcal{P}^{K}_{\text{off,regret}})$, $(\mathcal{P}^{\Gamma,K}_{\text{off,risk}})$, or $(\mathcal{P}^{\Gamma,K}_{\text{off,regret}})$; adapted from Vayanos et al. (2019), Subramanyam et al. (2017).

Inputs: Instance of Problem $(\mathcal{P}^{K}_{\text{off,risk}})$, $(\mathcal{P}^{K}_{\text{off,regret}})$, $(\mathcal{P}^{\Gamma,K}_{\text{off,risk}})$, or $(\mathcal{P}^{\Gamma,K}_{\text{off,regret}})$;

Output: Conservative (suboptimal) set of K queries $\{\boldsymbol{\iota}^{\kappa}\}_{\kappa\in\mathcal{K}}$ to ask the user;

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

  if $\kappa = 1$ **then**

    Solve the MBLP reformulation of Problem $(\mathcal{P}^{1}_{\text{off,risk}})$, $(\mathcal{P}^{1}_{\text{off,regret}})$, $(\mathcal{P}^{\Gamma,1}_{\text{off,risk}})$, or $(\mathcal{P}^{\Gamma,1}_{\text{off,regret}})$;

    Let $\boldsymbol{\iota}^{\star,1}$ denote an optimal query;

  else

    Solve the MBLP reformulation of Problem $(\mathcal{P}^{\kappa}_{\text{off,risk}})$, $(\mathcal{P}^{\kappa}_{\text{off,regret}})$, $(\mathcal{P}^{\Gamma,\kappa}_{\text{off,risk}})$, or $(\mathcal{P}^{\Gamma,\kappa}_{\text{off,regret}})$ with the added constraints that $\boldsymbol{\iota}^{k} = \boldsymbol{\iota}^{\star,k}$ for all $k \in \{1,\ldots,\kappa-1\}$;

    Let $\{\boldsymbol{\iota}^{\star,k}\}_{k=1}^{\kappa}$ denote an optimal query;

  end

end

Result: Return $\{\boldsymbol{\iota}^{\star,\kappa}\}_{\kappa\in\mathcal{K}}$.