

Robust Active Preference Elicitation*

Phebe Vayanos[†], Yingxiao Ye, Duncan C. McElfresh, John P. Dickerson, Eric Rice

[†]University of Southern California, Center for Artificial Intelligence in Society, phebe.vayanos@usc.edu

We study the problem of strategically eliciting the preferences of a decision-maker through a moderate number of pairwise comparison queries with the goal of making them a high quality recommendation for a specific decision-making problem. We are particularly motivated by applications in high stakes domains, such as when choosing a policy for allocating scarce resources to satisfy basic human needs (e.g., kidneys for transplantation or housing for those experiencing homelessness) where a consequential recommendation needs to be made from the (partially) elicited preferences. We model uncertainty in the preferences as being set based and 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 robust optimization formulations of these problems which integrate the preference elicitation and recommendation phases with aim to either maximize worst-case utility or minimize worst-case regret, and study their complexity. For the offline case, where active preference elicitation takes the form of a *two and half stage* robust optimization problem with decision-dependent information discovery, we provide 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 active preference learning takes the form of a *multi-stage* robust optimization problem with decision-dependent information discovery, we propose a conservative solution approach. Numerical studies on synthetic data demonstrate that our methods outperform state-of-the-art approaches from the literature in terms of worst-case rank, regret, and utility.

We showcase how our methodology can be used to assist a homeless services agency in choosing a policy for allocating scarce housing resources of different types to people experiencing homelessness. We use historical data to generate candidate counterfactual housing allocation policies that differ in their fairness, efficiency, and interpretability characteristics and elicit the agency’s values over these attributes.

Key words: decision-dependent information discovery, eliciting value judgements, high stakes problems.

1. Introduction

Automated recommendation and decision-support systems are routinely used in many aspects of our lives, e.g., to select adds to display (Li et al. 2010), to suggest products to purchase (Gomez-Uribe and Hunt 2016), or a route to take (Dai et al. 2015). Increasingly, they are used to make recommendations about consequential decisions in high stakes settings, e.g., to suggest a treatment plan to follow (Blasiak et al. 2020) or to prioritize transplant patients for deceased donor kidneys (Bertsimas et al. 2013). Yet, many of these decisions are multifaceted, requiring hard trade-offs to be made. For example, when choosing a treatment, one needs to balance its effectiveness (e.g., reduce pain) with potential risks (e.g., side effects), see Johnson et al. (2013).

Thus, to be able to make a recommendation likely to be adopted, the decision-support system needs to *elicit* the preferences of its user. Preference elicitation is achieved through either revealed

* Previous title: *Active preference elicitation via adjustable robust optimization*

preference data (Samuelson 1948) or stated preference surveys (Kroes and Sheldon 1988). The most common surveys are *conjoint studies* (Green and Srinivasan 1978, 1990) which model choices as vectors of attributes (as in multi-attribute utility theory (Dyer et al. 1992)) and characterize preferences in terms of a vector whose elements indicate how much the user values each attribute. The survey consists of several queries that ask the user to either select their preferred item among two alternatives (choice-based conjoint) or to state the strength of their preference between two items (metric-paired comparisons). The former is often preferred as it yields fewer response errors and surveys typically involve a moderate number of queries (~ 10) to avoid tiring the user.

In this paper, we study the problem of eliciting the preferences of a user using a moderate number of pairwise comparison (choice-based conjoint) queries with the goal of making them a high quality recommendation for a *specific* decision-making problem. We call this problem, where elicitation and recommendation phases are inextricably linked, *active preference elicitation*. We are particularly motivated by applications in high stakes low resource domains where revealed preference data is not available, where the number of attributes that characterize a recommendation is often large, and where a consequential recommendation needs to be made from the (partially) elicited preferences. We consider both an *offline* elicitation setting, where all queries are selected in advance, and an *online* elicitation setting, where the choice of query can be adapted to previous questions and answers. The former is relevant when working in low resource settings without access to a computer or when seeking to compare answers across users. The latter should be used whenever feasible to yield better recommendations. Given the critical nature of the problems we focus on, we aim to make *robust recommendations* that either maximize worst-case utility or minimize worst-case regret accounting for information incompleteness, inconsistencies in the responses, and potential model misspecifications. We aim to identify queries that are as informative as possible for the recommendation problem even when answers to queries are “adversarial” in the sense that they provide as little information as possible to guide the choice of recommendation. This point of view ensures our system will make recommendations with guaranteed performance. Technically, this problem can be cast as a robust sequential decision-making problem where the answers to the queries are uncertain and will only be revealed if the corresponding query is asked; such problems where information is endogenous are often said to present *decision-dependent information discovery* (DDID). We argue that active preference elicitation in this context is highly relevant to practical problems and requires the development of a new methodological framework.

A concrete motivation for our research is the plight of people experiencing homelessness and the ethical dilemma of deciding how to prioritize them for scarce housing resources. For example, in Los Angeles there are over 66,000 persons experiencing homelessness on any given night and only 22,769 housing units became available over the past year. This shortage makes it hard to decide

on an allocation policy as moral trade-offs need to be made between efficiency (e.g., having the least possible returns to homelessness) and equity (e.g., giving people resources commensurate to their needs). These challenges are further complicated by preexisting disparities. For example, in LA, Black people are four times more represented among those experiencing homelessness than in the general population (LAHSA 2018). Some agencies may prefer allocation rules that do not exacerbate inequalities, while others may prefer giving people equal chances. From our discussions with policy-makers at the Los Angeles Homeless Services Authority (LAHSA), the entity in charge of allocating housing resources in LA, there is consensus that the current matching mechanism is inadequate as it does not align with particular policy objectives, resulting for example in Black people returning to homelessness at almost twice the rate of White individuals (LAHSA 2018). While selecting a policy among two candidates is feasible, articulating a full preference profile is very challenging. At the same time, the ethical implications of housing allocation require it to closely reflect the interests of its users. Systems that do not match the needs of those they seek to help may deter them from seeking access. Systems that do not serve the interests of their users will result in unintended consequences and may get defunded, exacerbating resource scarcity and causing additional strain. The role of policy-makers in this context is to find ways to finely balance the often conflicting interests of those they serve in a way that aligns with their own values.

In recent years, several artificial intelligence driven tools have been proposed to design policies for matching homeless persons to housing (Azizi et al. 2018, Kube et al. 2019). However, these methods assume that the preferences of policy-makers are perfectly known. For example, Azizi et al. (2018) proposed a method for designing interpretable policies that minimize returns to homelessness while being, in some predefined sense, fair. To the best of our knowledge, no tool exists for eliciting the value judgements of policy-makers in housing allocation agencies to design policies that align with them. In fact, in private communications, policy-makers at LAHSA have described this as a recurring challenge. This is not surprising considering the many difficulties involved in identifying queries whose responses will be most useful in informing the ultimate choice of policy. First, selecting queries optimally requires solving a sequential decision-making problem with decision-dependent information discovery and combinatorial number of uncertain response scenarios. Second, it requires explicitly accounting for possible inconsistencies in the responses while also allowing for potential model misspecifications.

The problem we study is not unique to the housing allocation system but is rather encountered broadly in the design of policy-making systems, in particular those that decide how to allocate scarce resources needed to satisfy basic needs. Another such system is the US kidney allocation system (KAS), which provides deceased donor kidneys for transplant to those suffering from end stage renal disease. As of 2019, there were approximately 95,000 patients on the national kidney

waiting list and only 23,401 transplants took place. At the same time, there are severe inequities in kidney disease, with Black people being five times more likely to be affected than Whites (USRDS 2019). Thus, when deciding how to prioritize patients, the Organ Procurement and Transplantation Network (OPTN) faces ethical challenges and potential bad implications from inadequate allocation that are not unlike those faced in the housing allocation system.

Our research objective is to select the best possible queries to ask a user faced with a specific decision-making problem, taking into account any prior information about their preferences. The choice of queries should result in robust recommendations with the best possible guaranteed performance no matter what the user preferences are, and accounting for information incompleteness, possible inconsistencies, and potential model misspecifications, making them suitable to deploy in high stakes settings—we refer to these as *robust queries*. For example, we aim to elicit the moral value judgements of the leadership of a housing allocation agency to be able to recommend a policy that best aligns with them. The idea of examining and comparing sets of policies by looking at their outcomes is common in policy-making settings, e.g., it is used routinely in the KAS.

The preference elicitation problem has been widely studied in the artificial intelligence (AI), operations research (OR), and marketing literatures, see Section 2. Several researchers have investigated elicitation strategies and preference models that can be used to make robust recommendations. Yet, this large body of work is not well suited to make recommendations in high stakes low resource settings. Indeed, existing methods focus on obtaining unbiased estimates of preferences or on reducing uncertainty around estimates, rather than on selecting queries that will help explicitly improve the quality of the robust recommendation. This implies they are likely to result in suboptimal recommendations that do not align well with the values of end users. Moreover, none of the preference elicitation methods that can be used to make robust recommendations apply to the offline setting. We consequently utilize techniques from modern adjustable robust optimization that will be able to provide us with optimality guarantees on the quality of the queries selected while also allowing us to effectively cope with model misspecifications, inconsistencies, and information incompleteness.

The problem of selecting comparison queries that will explicitly maximize the quality of the downstream recommendation is a robust optimization problem with DDID (Vayanos et al. 2011, 2019). In particular, active offline (resp. online) elicitation can be cast as a *two and a half stage* (resp. *multi-stage*) problem with DDID. Our framework does not postulate distributions for uncertain parameters and instead models stochasticity by means of optimization variables that lie in uncertainty sets. This approach builds upon the line of work on polyhedral methods in conjoint analysis, see Section 2, but provides a different means of modelling inconsistencies that better aligns with popular choice models from the literature. From a modeling standpoint, the key feature that sets us apart from existing work is that our choice of queries is precisely guided by the structure

of the recommendation problem. From a solution perspective, this complicates significantly the problem of identifying good queries.

Indeed, when explicitly framed as a multi-stage robust optimization problem with DDID, the preference elicitation problem is particularly challenging: it involves binary adaptive decision variables (which also control the time of information discovery), discrete uncertainty (the responses to each queries) which may also be misspecified, an exponential number of response scenarios, a decision-dependent uncertainty set that involves “non-equalities” (to guarantee that responses are aligned with some utility model), and a decision-dependent uncertainty set that involves strict inequalities (to model strict preferences). This type of problem has not been previously studied in the literature on robust optimization, nor in the very recent work on robust optimization with DDID, see Section 2, and we show that it is generally NP-hard. By combining techniques from polyhedral theory with methods from robust optimization, we demonstrate that the offline active preference elicitation problems that maximize utility (with arbitrary recommendation set) or minimize regret (with enumerable recommendation set) can be reformulated equivalently as mixed-binary linear optimization problems (MBLP) involving a number of decision variables and constraints that are exponential in the number of queries and propose novel column-and-constraint generation approaches to solve them. For the online setting, we propose conservative solution approaches that solve a series of offline problems in a rolling horizon fashion.

We demonstrate the performance of our approaches in terms of quality of the recommendation and solution times by conducting extensive numerical studies using simulated data for realistic problem sizes. We also illustrate that changing the recommendation problem from max-min utility to min-max regret impacts the optimal choice of queries to be made, thereby showing the importance of coupling elicitation and recommendation phases.

We showcase how our methods can be used to assist policy-makers at agencies such as LAHSA in choosing a policy for allocating housing resources to people experiencing homelessness. We use historical data to generate candidate counterfactual allocation policies that differ in their fairness, efficiency, and interpretability characteristics and elicit stakeholder values over these attributes.

Our work contributes to the following literature streams. First, to preference elicitation by providing a framework for generating provably optimal robust queries for specific recommendation problems accounting for information incompleteness, model misspecifications, and response inconsistencies, making it suitable to deploy in high stakes settings. Second, to robust optimization with DDID by studying a model with discrete uncertain parameters, exponential number of contingencies, and uncertainty sets involving “non-equalities” and strict-inequalities. Finally, to the OR literature by developing the first framework for designing policies for allocating housing to those experiencing homelessness whose characteristics align with human value judgements.

Notation. Throughout this paper, vectors (matrices) are denoted by boldface lowercase (uppercase) letters. The k th element of a vector $\mathbf{x} \in \mathbb{R}^n$ ($k \leq n$) is denoted by \mathbf{x}_k . We let \mathbf{e} (resp. \mathbf{e}_i) denote the vector of all ones (resp. the i th 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.

Structure of the Paper. The remainder of the paper is organized as follows. Section 2 introduces related literature. Section 3 formalizes the model of our recommendation system and the preference model that underlies our approach. Sections 4 and 5 study the max-min utility and min-max regret active preference elicitation problems, respectively. Section 6 applies our framework to the problem of designing housing allocation policies that align with human value judgements and Section 7 concludes. The proofs of all statements can be found in the Electronic Companion to the paper.

2. Review of Related Work

Our work lies at the intersection of three literature streams: conjoint analysis in marketing, preference elicitation in AI, and stochastic and robust optimization with DDID in OR. It is also connected to preference robust optimization, multi-armed bandits, recommender systems, assortment planning, and dynamic pricing.

Conjoint Analysis. Most of the literature on conjoint analysis is focused on nonadaptive (offline) elicitation, relies on probability models for learning from discrete choices (McFadden 1974, 2001, Stern 1990), and builds on the field of experimental design. Queries are chosen to maximize a proxy of “precision” of the parameter estimates produced from the answers to the queries, see e.g., Arora and Huber (2001), Huber and Zwerina (1996), Kuhfeld et al. (1994), Kuhfeld (2010).

Several works consider adaptive (online) elicitation and model uncertainty in the utility by means of a polyhedron. This set, that we refer to as *uncertainty set* following robust optimization terminology, is progressively “shrunk” as answers to queries are observed. Toubia et al. (2003) considers metric paired-comparisons while Toubia et al. (2004, 2007) and Bertsimas and O’Hair (2013) focus on choice-based conjoint queries. These approaches leverage optimization techniques to select queries *one at a time* (greedily) in a way that will reduce the size of the resulting uncertainty set as much as possible. Inconsistencies are handled either by considering distributions supported by mixtures of polyhedra (Toubia et al. 2007) or by constraining the number of inconsistencies (Bertsimas and O’Hair 2013) and recommendations using the analytic center of the uncertainty set (Toubia et al. 2003, 2004, 2007) or by maximizing worst-case utility (Bertsimas and O’Hair 2013).

Modeling uncertainty using uncertainty sets as in the *polyhedral methods* is particularly attractive when seeking robust solutions—it circumvents the need for estimating distributions and for doing hard Bayesian updates. We thus build upon this line of research and model uncertainty using a polyhedron. However, our work has many differences. First, none of the polyhedral methods apply to the offline setting. Second, our framework explicitly models and solves the elicitation problem as a sequential decision-making problem. Third, our model of inconsistencies is different: it restricts errors in the responses to lie in a polyhedral set. Compared with the approach of Toubia et al. (2007), it does not result in a combinatorial explosion in the number of polyhedra. Relative to Bertsimas and O’Hair (2013), it ensures that response errors that are unlikely will be excluded, resulting in a more realistic, less conservative model. These differences ensure that our framework will result in queries with better performance in terms of worst-case utility or worst-case regret.

Recently, Sauré and Vielma (2019) proposed a new approach to adaptive choice-based conjoint analysis that is based on Bayesian updating but that preserves some of the attractive interpretability properties of the polyhedral method. As this approach decouples elicitation and recommendation while also taking a probabilistic perspective, it is not well suited to the problems we study.

Preference Elicitation in AI. Online preference elicitation has received considerable attention in AI. Some of this work takes a Bayesian perspective and relies on heuristics to identify queries (Chajewska et al. 2000, Boutilier 2002, Brochu et al. 2007, Zhao et al. 2018). Others model uncertainty using polyhedra and focus on gamble queries (Wang and Boutilier 2003, Boutilier et al. 2006). Utility is usually elicited using heuristics and recommendations are made using the min-max regret criterion (Viappiani and Kroer 2013). In our work, we also investigate recommendations that minimize worst-case regret but focus on choice-based conjoint queries and study also offline elicitation. The AI community has investigated many applications of preference elicitation, see e.g., Sandholm and Boutilier (2006), Braziunas and Boutilier (2010), Freedman et al. (2018), Benade et al. (2017).

Stochastic and Robust Optimization with DDID. Stochastic programming (SP) with DDID dates back to the works of Jonsbråten et al. (1998) and Jonsbråten (1998). Most of this literature assumes that the uncertain parameters are discretely distributed and models uncertainty by means of a finite scenario tree whose structure depends on the binary decisions that determine the time of information discovery (Goel and Grossman 2004, 2006, Goel et al. 2006, Colvin and Maravelias 2010, Gupta and Grossmann 2011). These methods do not apply to our context which involves also real valued uncertain parameters and we wish to optimize decisions robustly.

Our work most closely relates to RO with DDID (Vayanos et al. 2011, 2019). 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) study a solution method based on the K -adaptability approximation. They also investigate an active preference elicitation problem in

the context of the US KAS which is however very different from the one we consider in the present paper. Indeed, queries in Vayanos et al. (2019) correspond to a single item and the user is asked to quantify “how much” they like the item (on a continuous scale from 0 to 1). Thus, uncertainty is real-valued and moderately sized whereas our problem presents an exponential search space and an exponential number of binary contingencies, requiring a fundamentally different solution approach. In addition, our solution to offline elicitation is exact, whereas their method is approximate.

The problems we study also relate to the works of Nohadani and Roy (2017) and Iancu and Trichakis (2018) which focus on robust monitoring problems. Nohadani and Roy (2017) consider problems with time-dependent uncertainty sets that capture increased uncertainty as information ages and observations “reset uncertainty”. Iancu and Trichakis (2018) consider a system with evolving state that can be stopped at any time by the decision-maker to yield a time dependent reward. The decision-maker can choose, in conjunction with the stopping policy, a limited number of monitoring times when to observe the state of the system. Contrary to these problems the state of the system does not change in our case, there is a large number of parameters that we can choose to observe, uncertainty is discrete, and there is a very large number of possible scenarios.

Preference Robust Optimization. There is very rich literature connecting stochastic and robust optimization and risk preferences (Dentcheva and Ruszczyński 2004, Natarajan et al. 2007, Brown and Sim 2009, Chen and Sim 2009, Brown et al. 2012, Lam et al. 2013). Our paper most closely relates to preference robust optimization. This literature assumes that partial information on preferences is available (e.g., through previous surveys) and focuses on making robust recommendations accounting for ambiguity in risk preferences while capturing sophisticated choice models (Armbruster and Delage 2015, Hu and Mehrotra 2015, Hu et al. 2018, Delage and Li 2018, Delage et al. 2018, Haskell et al. 2018). Differing from this work we study the problem of *selecting queries* and focus on learning trade-offs between attributes rather attitudes towards risk.

Multi-Armed Bandits. Multi-armed bandits (MAB) offer a framework for modeling problems where information is acquired over time as a by-product of decisions (Jones and Gittins 1989). Extant models account for capacity constraints (Wu et al. 2015), contextual information (Lu et al. 2010, Badanidiyuru et al. 2014), and provide robust formulations (Auer et al. 1995, 2003, Syrgkanis et al. 2016a,b, Lykouris et al. 2018). The key feature that differentiates our setting is that we do not obtain a reward as a byproduct of the decisions and that the number of arms is exponential if we model our problem from this lens. Thus MAB are not suited to tackle our problem.

Recommender Systems. Research on recommender systems aims to help users navigate large product spaces and has been motivated from web applications, e.g., to generate playlists for video and music services (Ladeira 2019), or content for social media platforms (Gupta et al. 2013). They aim to *filter* content to provide users with an experience that aligns with their preferences while still

being rich/diverse enough, see e.g., McSherry (2002), Reilly et al. (2005). These methods generally rely on a lot more data than we can collect and are not well suited for high stakes settings.

Assortment Planning. Static assortment optimization with known choice behavior has been an active area of research (Van Ryzin and Mahajan 1999, Mahajan and Van Ryzin 2001). Dynamic assortment optimization which adaptively learns the customers’ choice behavior is receiving increasing attention. Most of this work does not consider the features of the items (Caro and Gallien 2007, Rusmevichientong et al. 2010, Sauré and Zeevi 2013, Agrawal et al. 2019) or focuses on the pricing problem (Roth et al. 2016, 2020). Our work most closely relates to assortment planning with contextual models (Cheung and Simchi-Levi 2017, Chen et al. 2020, Oh and Iyengar 2019). Chen et al. (2019) consider contextual information while also taking a robust viewpoint. However, their work is based on revealed preferences and focused on multiple customers.

Dynamic Pricing. The literature on dynamic pricing with learning studies pricing in settings where the demand function is unknown and can be learned as a byproduct of decisions (Kleinberg and Leighton 2003). A large body of literature has studied this topic including both parametric (Araman and Caldentey 2009, Broder and Rusmevichientong 2012, Harrison et al. 2012, Chen and Farias 2013, den Boer and Zwart 2013, Besbes and Zeevi 2012) and non-parametric (Besbes and Zeevi 2009, Harrison et al. 2013) approaches. The literature also includes models that use a robust optimization approach (Bertsimas and Vayanos 2017, Cohen et al. 2020). Our work is different in that we need to learn preferences from pairwise comparison queries and stated preference surveys.

3. Model

In this section, we define the building blocks of our framework that will be used to compute optimal active preference elicitation strategies in Sections 4 and 5.

3.1. Items, Query and Recommendation Sets, and User Preference Model

Items. We assume that, when choosing one product (e.g., policy) over another, a user is basing their decision on the attributes of the options, as in conjoint analysis. Thus, we characterize each item \mathbf{x} by its J attributes and model it as a point in a J -dimensional Cartesian space, i.e., $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_J)$, where $\mathbf{x}_j \in \mathbb{R}$, $j = 1, \dots, J$, denotes the j th attribute of item $\mathbf{x} \in \mathbb{R}^J$. In this work, we are motivated by settings where J is high-dimensional (in the order of 10 or 20 for example). We denote the universe of all item configurations that are possible to produce by $\mathcal{X} \subseteq \mathbb{R}^J$.

Query 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 assume that query set \mathcal{Q} has finite cardinality $I \geq 2$ and index items in this set by $i \in \mathcal{I} := \{1, \dots, I\}$. We denote the i th item by $\mathbf{x}^i \in \mathbb{R}^J$, i.e., $\mathcal{Q} := \{\mathbf{x}^i \mid i \in \mathcal{I}\}$. This assumption is very common and also natural—e.g., in policy-making settings such as during the OPTN Kidney Transplantation Committee Meetings, there is usually a finite number of candidate policies that are being compared.

Recommendation Set. We denote by $\mathcal{R} \subseteq \mathcal{X}$ the recommendation set, i.e., the set of items from which the system can draw to make a recommendation or equivalently the feasible set of the optimization problem that the decision-maker is interested in solving. In general, \mathcal{Q} and \mathcal{R} need not coincide. For example, we may be able to design a new item, that is not currently available or we may be unable to recommend a product that is now out-of-stock. In our work, the set \mathcal{R} may have a very general structure. For example, it may take the form of the feasible set of a shortest path problem, a knapsack problem, an assignment problem, or a minimum spanning tree problem.

User Preference Model. We assume that the user's preferences can be represented with a linear utility function $u: \mathcal{X} \rightarrow \mathbb{R}$ defined through $u(\mathbf{x}) := \mathbf{u}^\top \mathbf{x} + \gamma(\mathbf{x})$, where $\mathbf{u} \in \mathbb{R}^J$ is a vector representing the relative importance they assign to product attributes and $\gamma(\mathbf{x})$ is an idiosyncratic shock to utility. We make the following assumption regarding the vector \mathbf{u} .

Assumption 1 (Polyhedral Uncertainty). *The vector \mathbf{u} is an element of the uncertainty set $\mathcal{U}^0 \subseteq \mathbb{R}^J$ defined through $\mathcal{U}^0 := \{\mathbf{u} \in \mathbb{R}^J \mid \mathbf{B}\mathbf{u} \geq \mathbf{b}\}$ for some matrix $\mathbf{B} \in \mathbb{R}^{M \times J}$ and vector $\mathbf{b} \in \mathbb{R}^M$. Moreover, \mathcal{U}^0 is non-empty, full-dimensional, and bounded.*

The assumption that u is linear and that prior information on \mathbf{u} can be encoded using linear inequalities is standard in the literature on preference elicitation and polyhedral methods to conjoint analysis, see Section 2. Our utility model is however different than the ones used in this line of work in that we explicitly assume that the utility function is affected by idiosyncratic shocks, in line with several popular choice models from the literature, see e.g., Greene (2012). This modelling choice will allow us to have a polyhedral model for the inconsistencies, see Section 3.2.

Assumption 1 can be relaxed to allow that \mathcal{U}^0 presents equality constraints and that it possesses an inner point (rather than being full-dimensional) at the cost of complicating our proofs.

Example 1. *If no prior information is available on \mathbf{u} , since the utility coefficients can be scaled by a constant without affecting the preference ordering over items, one may use $\mathcal{U}^0 = [-1, 1]^J$, see e.g., Bertsimas and O'Hair (2013).*

Example 2. *Alternatively, one may set $\mathcal{U}^0 = \{\mathbf{u} \in \mathbb{R}_+^J : \mathbf{e}^\top \mathbf{u} = 1\}$ as in Toubia et al. (2003). In that case, the coefficients \mathbf{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 3. *If item $\bar{\mathbf{x}}$ is a benchmark, one may normalize utilities relative to that item, by letting $\mathcal{U}^0 = \{\mathbf{u} \in [-1, 1]^J : \mathbf{u}^\top \bar{\mathbf{x}} = 1\}$.*

3.2. Elicitation through Comparison Queries

Before recommending an item from the set \mathcal{R} , the recommender system has the opportunity to make K queries, indexed in the set $\mathcal{K} := \{1, \dots, K\}$, to the user. These queries may enable the system to gain information about \mathbf{u} , thereby improving the quality of the recommendation. Each query takes the form of a comparison between two items in the query set, thus being an element of $\mathcal{C} := \{(i, i') \mid i \in \mathcal{I}, i' \in \mathcal{I}, i < i'\}$. We let $\boldsymbol{\iota}^\kappa := (\boldsymbol{\iota}_1^\kappa, \boldsymbol{\iota}_2^\kappa) \in \mathcal{C}$ denote the κ th 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 κ th query, respectively: query κ asks the user to compare $\mathbf{x}^{\boldsymbol{\iota}_1^\kappa}$ and $\mathbf{x}^{\boldsymbol{\iota}_2^\kappa}$. The user must choose one of three possible answers in response to query κ : (a) ‘‘I strictly prefer $\mathbf{x}^{\boldsymbol{\iota}_1^\kappa}$ to $\mathbf{x}^{\boldsymbol{\iota}_2^\kappa}$ ’’ (i.e., $\mathbf{x}^{\boldsymbol{\iota}_1^\kappa} \succ \mathbf{x}^{\boldsymbol{\iota}_2^\kappa}$); (b) ‘‘I am indifferent between $\mathbf{x}^{\boldsymbol{\iota}_1^\kappa}$ and $\mathbf{x}^{\boldsymbol{\iota}_2^\kappa}$ ’’ (i.e., $\mathbf{x}^{\boldsymbol{\iota}_1^\kappa} \sim \mathbf{x}^{\boldsymbol{\iota}_2^\kappa}$); or (c) ‘‘I strictly prefer $\mathbf{x}^{\boldsymbol{\iota}_2^\kappa}$ to $\mathbf{x}^{\boldsymbol{\iota}_1^\kappa}$ ’’ (i.e., $\mathbf{x}^{\boldsymbol{\iota}_1^\kappa} \prec \mathbf{x}^{\boldsymbol{\iota}_2^\kappa}$). We associate each possible answer to query κ to a response scenario $\mathbf{s}_\kappa \in \mathcal{S} := \{-1, 0, 1\}$ such that

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

The information obtained on \mathbf{u} depends on the answer to query κ , i.e., on the response scenario. Since the utility of any given item is affected by idiosyncratic noise, the answers given by the user may be *inconsistent* in the sense that $\mathbf{x} \succ \mathbf{y}$ and $u(\mathbf{x}) \leq u(\mathbf{y})$; $\mathbf{x} \prec \mathbf{y}$ and $u(\mathbf{x}) \geq u(\mathbf{y})$; or $\mathbf{x} \sim \mathbf{y}$ and $u(\mathbf{x}) \neq u(\mathbf{y})$. Naturally, such inconsistencies may not occur if the user is perfectly rational, see e.g., Arrow (1963). Motivated by our utility model, see Section 3.1, we impose a different linear constraint on \mathbf{u} depending on the response: (a) If $\mathbf{s}_\kappa = 1$, then $\mathbf{u}^\top \mathbf{x}^{\boldsymbol{\iota}_1^\kappa} - \mathbf{u}^\top \mathbf{x}^{\boldsymbol{\iota}_2^\kappa} > -\epsilon_\kappa$; (b) If $\mathbf{s}_\kappa = 0$, then $|\mathbf{u}^\top \mathbf{x}^{\boldsymbol{\iota}_1^\kappa} - \mathbf{u}^\top \mathbf{x}^{\boldsymbol{\iota}_2^\kappa}| \leq \epsilon_\kappa$; and (c) If $\mathbf{s}_\kappa = -1$, then $\mathbf{u}^\top \mathbf{x}^{\boldsymbol{\iota}_1^\kappa} - \mathbf{u}^\top \mathbf{x}^{\boldsymbol{\iota}_2^\kappa} < \epsilon_\kappa$, where $\boldsymbol{\epsilon} := \{\epsilon_\kappa\}_{\kappa \in \mathcal{K}}$ is a vector of inconsistencies. We make the following assumption about $\boldsymbol{\epsilon}$.

Assumption 2 (Polyhedral Inconsistencies). *The user inconsistencies across K queries are assumed to lie in the uncertainty set $\mathcal{E}_\Gamma := \{\boldsymbol{\epsilon} \in \mathbb{R}_+^K : \sum_{\kappa \in \mathcal{K}} \epsilon_\kappa \leq \Gamma\}$ for some (known) parameter $\Gamma \in \mathbb{R}_+$, referred to as the budget of uncertainty.*

Example 4 (Choice of Γ). *Standard approaches in RO can be used to choose Γ (Ben-Tal et al. 2009, Bertsimas et al. 2010). For example, if the idiosyncratic shocks in utility are assumed to be normal with mean zero and standard deviation σ , then $\sum_{\kappa \in \mathcal{K}} \epsilon_\kappa$ is normal with standard deviation $2\sigma\sqrt{K}$ and setting $\Gamma = 2\sigma\sqrt{K}\text{erf}^{-1}(2p - 1)$ ensures that $\boldsymbol{\epsilon} \in \mathcal{E}_\Gamma$ with probability p .*

Modeling inconsistencies as lying in a polyhedron sets us apart from the literature which either uses mixtures of polyhedra for the support of \mathbf{u} (Toubia et al. 2007), or integer programming to bound the number of incorrect responses and big- M constants to toggle the sign of the constraints added to \mathcal{U}_0 (Bertsimas and O’Hair 2013). In particular, our modeling paradigm is more interpretable

than the approach of Toubia et al. (2007). In addition, it is less conservative than the framework of Bertsimas and O’Hair (2013) since our model will only allow inconsistent responses when the user “marginally” prefers one item over the other. Moreover, modeling inconsistencies in this way allows us to use duality theory to solve our problems in Sections 4 and 5.

We collect the answers to each query in the vector $\mathbf{s} := \{\mathbf{s}_\kappa\}_{\kappa \in \mathcal{K}}$ and accordingly let $\boldsymbol{\nu} := \{(\boldsymbol{\nu}_1^\kappa, \boldsymbol{\nu}_2^\kappa)\}_{\kappa \in \mathcal{K}}$. After all K queries have been made and the responses to the queries are observed, we can update \mathcal{U}^0 as:

$$\mathcal{U}_\Gamma(\boldsymbol{\nu}, \mathbf{s}) := \left\{ \mathbf{u} \in \mathcal{U}^0 : \exists \boldsymbol{\epsilon} \in \mathcal{E}_\Gamma \text{ such that } \begin{array}{ll} \mathbf{u}^\top (\mathbf{x}^{\boldsymbol{\nu}_1^\kappa} - \mathbf{x}^{\boldsymbol{\nu}_2^\kappa}) > -\epsilon_\kappa & \forall \kappa \in \mathcal{K} : \mathbf{s}_\kappa = 1 \\ |\mathbf{u}^\top (\mathbf{x}^{\boldsymbol{\nu}_1^\kappa} - \mathbf{x}^{\boldsymbol{\nu}_2^\kappa})| \leq \epsilon_\kappa & \forall \kappa \in \mathcal{K} : \mathbf{s}_\kappa = 0 \\ \mathbf{u}^\top (\mathbf{x}^{\boldsymbol{\nu}_1^\kappa} - \mathbf{x}^{\boldsymbol{\nu}_2^\kappa}) < \epsilon_\kappa & \forall \kappa \in \mathcal{K} : \mathbf{s}_\kappa = -1 \end{array} \right\}.$$

We emphasize that after the K comparisons, the set of possible values for \mathbf{u} depends both on $\boldsymbol{\nu}$, the queries made, and on \mathbf{s} , the answers given by the user. This idea is illustrated on Figure EC.1.

3.3. Robust Recommendations

After K queries have been made and the answers to these queries have been observed, the recommendation system needs to select an item from the set \mathcal{R} to recommend albeit the coefficients \mathbf{u} are only known to belong to the uncertainty set $\mathcal{U}_\Gamma(\boldsymbol{\nu}, \mathbf{s})$. We thus propose to provide recommendations that are robust to all possible realization of \mathbf{u} in the set $\mathcal{U}_\Gamma(\boldsymbol{\nu}, \mathbf{s})$ and investigate two notions of robustness that are popular in the literature: max-min utility and min-max regret.

Maximizing Worst-Case Utility. Given uncertainty in the utility function coefficients, it is natural to seek recommendations that will maximize worst-case utility. Mathematically, given the sequences, $\boldsymbol{\nu}$ and \mathbf{s} , the recommender system offers an item \mathbf{x} that solves the problem

$$\underset{\mathbf{x} \in \mathcal{R}}{\text{maximize}} \quad \min_{\mathbf{u} \in \mathcal{U}_\Gamma(\boldsymbol{\nu}, \mathbf{s})} \quad \mathbf{u}^\top \mathbf{x}. \quad (\mathcal{RP}_u)$$

Minimizing Worst-Case Regret. An alternative is to make recommendations according to the worst-case absolute regret criterion which compares 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 (which maximizes worst-case utility) and is thus attractive in practical applications. According to this criterion and given the sequences $\boldsymbol{\nu}$ and \mathbf{s} , the recommender system offers the item that solves the problem

$$\underset{\mathbf{x} \in \mathcal{R}}{\text{minimize}} \quad \max_{\mathbf{u} \in \mathcal{U}_\Gamma(\boldsymbol{\nu}, \mathbf{s})} \quad \left\{ \max_{\mathbf{x}' \in \mathcal{R}} \mathbf{u}^\top \mathbf{x}' - \mathbf{u}^\top \mathbf{x} \right\}. \quad (\mathcal{RP}_r)$$

We note that for some choices of $\boldsymbol{\iota}$ and \boldsymbol{s} , the set $\mathcal{U}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s})$ may be open and even empty and emphasize that the sets of optimal solutions to (\mathcal{RP}_u) and (\mathcal{RP}_r) depend on both $\boldsymbol{\iota}$ and \boldsymbol{s} . While the system has no control over the responses \boldsymbol{s} it *can* select the queries $\boldsymbol{\iota} \in \mathcal{C}^K$ to improve the quality of the recommendation. This observation will motivate our formulations in Sections 4 and 5.

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

In this section, we study the offline and online active preference elicitation problems under the max-min utility decision criterion and evaluate their performance on synthetic data.

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

In the offline problem with the max-min utility decision criterion, all K comparisons are precommitted *in advance*, before any agent responses are observed and the goal is to select robust queries that will result in highest worst-case utility of the recommendation. Offline elicitation is relevant when interactions with the user do not involve a computer (e.g., paper-based questionnaires) or if all users must be presented the same queries (e.g., in a controlled study). It is also a building block for solving the online preference elicitation problem, see Section 4.2.

4.1.1. Problem Formulation & Complexity Analysis In the offline problem with the max-min utility decision criterion, 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}^* from \mathcal{U}^0 , responds to the queries truthfully (albeit potentially inconsistently, see Section 3.2) by selecting answers \boldsymbol{s}_κ to each query $\boldsymbol{\iota}^\kappa$ in a way that complies with their utility vector \boldsymbol{u}^* and some choice of $\boldsymbol{\epsilon} \in \mathcal{E}_\Gamma$. Note that the utility vector \boldsymbol{u}^* is not observable to the recommender system; in fact, the user themselves is not aware of their vector \boldsymbol{u}^* (else, they would directly share it). Once the answers to the queries are observed, the recommender system can certify that $\boldsymbol{u}^* \in \mathcal{U}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s})$ and solves problem (\mathcal{RP}_u) .

Mathematically, the offline problem with the max-min utility decision criterion is expressible as the two-stage robust optimization problem with decision-dependent information discovery

$$\underset{\boldsymbol{\iota} \in \mathcal{C}^K}{\text{maximize}} \quad \underset{\boldsymbol{s} \in \mathcal{S}_\Gamma(\boldsymbol{\iota})}{\min} \quad \underset{\boldsymbol{x} \in \mathcal{R}}{\max} \quad \underset{\boldsymbol{u} \in \mathcal{U}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s})}{\min} \quad \boldsymbol{u}^\top \boldsymbol{x}, \quad (\mathcal{QP}_u^K)$$

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 $\boldsymbol{\epsilon} \in \mathcal{E}_\Gamma$. Indeed, the set $\mathcal{U}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s})$ is empty if and only if the answers given by the user are not consistent with any $\boldsymbol{u} \in \mathcal{U}^0$ and no vector $\boldsymbol{\epsilon} \in \mathcal{E}_\Gamma$ can resolve the inconsistencies. Since \mathcal{U}^0 is non-empty by Assumption 1, $\mathcal{S}_\Gamma(\boldsymbol{\iota}) \neq \emptyset$ for all $\boldsymbol{\iota} \in \mathcal{C}^K$. Problem (\mathcal{QP}_u^K) is a “two-and-a-half” stage dynamic robust optimization problem with decision-dependent uncertainty set.

Problem (\mathcal{QP}_u^K) is difficult to solve for many reasons. First, it is a max-min-max-min problem. Second, the uncertainty sets $\mathcal{S}_\Gamma(\boldsymbol{\iota})$ and $\mathcal{U}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s})$ for the first and second decision-stages are decision-dependent. Third, the uncertainty set $\mathcal{U}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s})$ depends upon the first stage uncertainty \boldsymbol{s} and is

also open, making it difficult to derive computational solution approaches. Fourth, the uncertainty set $\mathcal{S}_\Gamma(\boldsymbol{\iota})$ involves “non-equalities”. The following lemma shows that Problem (\mathcal{QP}_u^K) can be considerably simplified by eliminating the dependence of $\mathcal{S}_\Gamma(\boldsymbol{\iota})$ on $\boldsymbol{\iota}$, by replacing the strict inequalities in the set $\mathcal{U}_\Gamma(\boldsymbol{\iota}, \mathbf{s})$ by their loose counterparts, and by eliminating indifferent responses from the response set.

Lemma 1. *Problem (\mathcal{QP}_u^K) is equivalent to*

$$\underset{\boldsymbol{\iota} \in \mathcal{C}^K}{\text{maximize}} \quad \underset{\mathbf{s} \in \tilde{\mathcal{S}}^K}{\min} \quad \underset{\mathbf{x} \in \mathcal{R}}{\max} \quad \underset{\mathbf{u} \in \tilde{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \mathbf{s})}{\min} \quad \mathbf{u}^\top \mathbf{x}, \quad (\widetilde{\mathcal{QP}}_u^K)$$

where $\tilde{\mathcal{S}} := \{-1, 1\}$ and

$$\tilde{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \mathbf{s}) := \left\{ \begin{array}{l} \mathbf{u} \in \mathcal{U}^0 : \exists \boldsymbol{\epsilon} \in \mathcal{E}_\Gamma \text{ such that} \\ \mathbf{u}^\top (\mathbf{x}^{\boldsymbol{\iota}_1^{\kappa}} - \mathbf{x}^{\boldsymbol{\iota}_2^{\kappa}}) \geq -\boldsymbol{\epsilon}_\kappa \quad \forall \kappa \in \mathcal{K} : \mathbf{s}_\kappa = 1 \\ \mathbf{u}^\top (\mathbf{x}^{\boldsymbol{\iota}_1^{\kappa}} - \mathbf{x}^{\boldsymbol{\iota}_2^{\kappa}}) \leq \boldsymbol{\epsilon}_\kappa \quad \forall \kappa \in \mathcal{K} : \mathbf{s}_\kappa = -1 \end{array} \right\},$$

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

The above result avoids us having to use “little- ϵ ” variables to convert the strict inequalities in the uncertainty set to loose inequalities as done e.g., by Bertsimas and O’Hair (2013) in their procedure to calculate the analytic center of the uncertainty set.

Next, we show that Problem $(\widetilde{\mathcal{QP}}_u^K)$ is \mathcal{NP} -hard when the recommendation set is discrete, which motivates the integer optimization based reformulation we provide in the next section.

Theorem 1. *The following claims about the complexity of Problem $(\widetilde{\mathcal{QP}}_u^K)$ hold true.*

- (a) *Suppose the recommendation set \mathcal{R} is convex. Then, Problem $(\widetilde{\mathcal{QP}}_u^K)$ is polynomially solvable. Moreover, it holds that Problems $(\widetilde{\mathcal{QP}}_u^K)$ and $(\widetilde{\mathcal{QP}}_u^0)$ 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 $(\widetilde{\mathcal{QP}}_u^K)$ and (\mathcal{RP}_u) coincide.*
- (b) *Evaluating the objective function of Problem $(\widetilde{\mathcal{QP}}_u^K)$ is \mathcal{NP} -hard even if the recommendation set \mathcal{R} consists of only two elements.*

The intuition behind claim (a) is akin to that of the minimax theorem (Von Neumann 1928). Indeed, under the conditions in the statement, \mathcal{R} and $\tilde{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \mathbf{s})$ are both convex and so observing $\mathbf{u} \in \tilde{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \mathbf{s})$ before choosing \mathbf{x} yields no benefit. Since this holds true for any choice of $\boldsymbol{\iota}$ and \mathbf{s} , there is no gain (in the worst-case) from asking queries. Motivated by the complexity results above, we henceforth assume the recommendation set in Problem $(\widetilde{\mathcal{QP}}_u^K)$ is discrete and MBLP representable. We provide a general, MBLP reformulation of Problem $(\widetilde{\mathcal{QP}}_u^K)$ in Section 4.1.2, a column-and-constraint generation approach in Section 4.1.3, and additional speed-up strategies in Section 4.1.4.

4.1.2. Exact MBLP Reformulation While $(\widetilde{\mathcal{QP}}_u^K)$ can in principle be solved by enumeration over $\boldsymbol{\iota}$ and \mathbf{s} , such an approach becomes computationally prohibitive as K grows. We thus propose instead to reformulate $(\widetilde{\mathcal{QP}}_u^K)$ as an MBLP by noting that we can exchange the order of the inner minimization and maximization problems, provided we index the recommendation by \mathbf{s} . This observation combined with classical robust optimization techniques yields the following lemma.

Lemma 2. *Problem $(\widetilde{\mathcal{QP}}_u^K)$ is equivalent to the following finite program*

$$\begin{array}{ll}
\text{maximize} & \tau \\
\text{subject to} & \tau \in \mathbb{R}, \boldsymbol{\iota} \in \mathcal{C}^K \\
& \left. \begin{array}{l}
\mathbf{x}^s \in \mathcal{R}, \boldsymbol{\alpha}^s \in \mathbb{R}_+^K, \boldsymbol{\beta}^s \in \mathbb{R}_+^M, \mu^s \in \mathbb{R}_- \\
\tau \leq \mathbf{b}^\top \boldsymbol{\beta}^s + \Gamma \mu^s \\
\sum_{\kappa \in \mathcal{K}} \mathbf{s}_\kappa (\mathbf{x}^{\iota_1^\kappa} - \mathbf{x}^{\iota_2^\kappa}) \boldsymbol{\alpha}_\kappa^s + \mathbf{B}^\top \boldsymbol{\beta}^s = \mathbf{x}^s \\
\boldsymbol{\alpha}^s + \mu^s \mathbf{e} \leq \mathbf{0}
\end{array} \right\} \forall \mathbf{s} \in \widetilde{\mathcal{S}}^K,
\end{array} \tag{1}$$

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

Problem (1) can be converted to an MBLP by encoding the choice of $\boldsymbol{\iota}^\kappa \in \mathcal{C}$, $\kappa \in \mathcal{K}$, using binary decisions $\mathbf{v}^\kappa \in \{0, 1\}^I$ and $\mathbf{w}^\kappa \in \{0, 1\}^I$ whose i th element is one if and only if item i is the first (resp. second) item in query κ . This result is summarized in the following theorem.

Theorem 2. *Problem (1) is equivalent to the following mixed-binary linear program*

$$\begin{array}{ll}
\text{maximize} & \tau \\
\text{subject to} & \tau \in \mathbb{R}, \mathbf{v}^\kappa, \mathbf{w}^\kappa \in \{0, 1\}^I, \kappa \in \mathcal{K} \\
& \boldsymbol{\alpha}^s \in \mathbb{R}_+^K, \boldsymbol{\beta}^s \in \mathbb{R}_+^M, \mu^s \in \mathbb{R}_-, \mathbf{x}^s \in \mathcal{R}, \mathbf{s} \in \widetilde{\mathcal{S}}^K \\
& \overline{\mathbf{v}}^{\mathbf{s}^\kappa}, \overline{\mathbf{w}}^{\mathbf{s}^\kappa} \in \mathbb{R}_+^I, \mathbf{s} \in \widetilde{\mathcal{S}}^K, \kappa \in \mathcal{K} \\
& \left. \begin{array}{l}
\tau \leq \mathbf{b}^\top \boldsymbol{\beta}^s + \Gamma \mu^s \\
\sum_{i \in \mathcal{I}} \mathbf{x}^i \sum_{\kappa \in \mathcal{K}} \mathbf{s}_\kappa (\overline{\mathbf{v}}_i^{\mathbf{s}^\kappa} - \overline{\mathbf{w}}_i^{\mathbf{s}^\kappa}) + \mathbf{B}^\top \boldsymbol{\beta}^s = \mathbf{x}^s \\
\boldsymbol{\alpha}^s + \mu^s \mathbf{e} \leq \mathbf{0}
\end{array} \right\} \forall \mathbf{s} \in \widetilde{\mathcal{S}}^K \\
& \mathbf{e}^\top \mathbf{v}^\kappa = 1, \mathbf{e}^\top \mathbf{w}^\kappa = 1, 1 - \mathbf{w}_i^\kappa \geq \sum_{i': i' \geq i} \mathbf{v}_{i'}^\kappa \quad \forall i \in \mathcal{I}, \forall \kappa \in \mathcal{K} \\
& \left. \begin{array}{l}
\overline{\mathbf{v}}^{\mathbf{s}^\kappa} \leq M \mathbf{v}^\kappa, \overline{\mathbf{v}}^{\mathbf{s}^\kappa} \leq \boldsymbol{\alpha}_\kappa^s \mathbf{e}, \overline{\mathbf{v}}^{\mathbf{s}^\kappa} \geq \boldsymbol{\alpha}_\kappa^s \mathbf{e} - M(\mathbf{e} - \mathbf{v}^\kappa) \\
\overline{\mathbf{w}}^{\mathbf{s}^\kappa} \leq M \mathbf{w}^\kappa, \overline{\mathbf{w}}^{\mathbf{s}^\kappa} \leq \boldsymbol{\alpha}_\kappa^s \mathbf{e}, \overline{\mathbf{w}}^{\mathbf{s}^\kappa} \geq \boldsymbol{\alpha}_\kappa^s \mathbf{e} - M(\mathbf{e} - \mathbf{w}^\kappa)
\end{array} \right\} \begin{array}{l} \forall \mathbf{s} \in \widetilde{\mathcal{S}}^K \\ \kappa \in \mathcal{K}, \end{array}
\end{array} \tag{2}$$

where M is a “big- M ” constant. Given an optimal solution $(\tau, \{\mathbf{v}^\kappa, \mathbf{w}^\kappa\}_{\kappa \in \mathcal{K}}, \{\boldsymbol{\alpha}^s, \boldsymbol{\beta}^s, \mathbf{x}^s, \boldsymbol{\mu}^s\}_{s \in \tilde{\mathcal{S}}^K}, \{\bar{\mathbf{v}}_i^{s\kappa}, \bar{\mathbf{w}}_i^{s\kappa}\}_{s \in \tilde{\mathcal{S}}^K, \kappa \in \mathcal{K}})$ to Problem (2), an optimal set of queries for (\mathcal{QP}_u^K) is

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

The first and third sets of constraints in Problem (2) exactly match those in (1). The second set of constraints is almost identical between the two problems with the exception that the choice of queries is now encoded using the \mathbf{v} and \mathbf{w} variables. The fourth set of constraints guarantees that $\boldsymbol{\iota}_1^\kappa < \boldsymbol{\iota}_2^\kappa$, i.e., $(\boldsymbol{\iota}_1^\kappa, \boldsymbol{\iota}_2^\kappa) \in \mathcal{C}$. The fifth and sixth sets of constraints linearize the products $\boldsymbol{\alpha}_\kappa^s \mathbf{v}^\kappa$ and $\boldsymbol{\alpha}_\kappa^s \mathbf{w}^\kappa$. Reformulation (2) enables us to solve (\mathcal{QP}_u^K) as an MBLP using off-the-shelf solvers. For fixed K , its size is polynomial in the size of the input. Yet, it is exponential in K . In the next section, we propose a decomposition strategy to speed-up solution as K grows to typical values in the order of 8 to 10. Additional strategies are discussed in Section 4.1.4.

4.1.3. Column-and-Constraint Generation The number of scenarios in Problem (2) is exponential in K . At the same time, we expect that only a moderate number of scenarios $s \in \tilde{\mathcal{S}}^K$ will be active in the epigraph constraint. Column-and-constraint generation (CCG) techniques are a popular approach for addressing problems that possess an exponential number of decision variables and constraints while presenting a decomposable structure, see e.g., Fischetti and Vigo (1997), Löbel (1998), Valério De Carvalho (1999), Mamer and McBride (2000), Feillet et al. (2010), Sadykov and Vanderbeck (2011), Muter et al. (2013), and Muter et al. (2018). In particular, CCG algorithms have been proposed for solving two-stage robust optimization problems by Zeng and Zhao (2013), Bertsimas and Shtern (2018), and Vayanos et al. (2019). However, these do not apply to problems with a mix of real and discrete uncertainty and/or to problems with two-and-a-half decision stages. We thus propose a new CCG algorithm to speed-up computation that applies when \mathcal{R} has finite cardinality (in which case (\mathcal{QP}_u^K) is \mathcal{NP} -hard, see Theorem 1). We describe our CCG procedure using Problem (1). Naturally, all problems solved would need to be converted to MBLPs using techniques similar to those employed in Theorem 2. We omit these to streamline presentation.

The key idea behind our algorithm is to decompose the problem into a relaxed main problem and a series of subproblems indexed by $s \in \tilde{\mathcal{S}}^K$. The main problem initially only involves a subset of the constraints (those indexed by $s \in \mathcal{S}' \subseteq \tilde{\mathcal{S}}^K$) and a *single auxiliary MBLP* is used to iteratively identify indices $s \in \tilde{\mathcal{S}}^K$ for which the solution to the main problem becomes infeasible when plugged into subproblem s . Constraints associated with infeasible subproblems are added to the main problem and the procedure continues until convergence. We now detail this approach.

We define the relaxed main problem parameterized by the index set \mathcal{S}' as

$$\begin{aligned}
& \text{maximize} && \tau \\
& \text{subject to} && \tau \in \mathbb{R}, \boldsymbol{\nu} \in \mathcal{C}^K \\
& && \left. \begin{aligned}
& \boldsymbol{\alpha}^s \in \mathbb{R}_+^K, \boldsymbol{\beta}^s \in \mathbb{R}_+^M, \mu^s \in \mathbb{R}_-, \mathbf{x}^s \in \mathcal{R} \\
& \tau \leq \mathbf{b}^\top \boldsymbol{\beta}^s + \Gamma \mu^s \\
& \sum_{\kappa \in \mathcal{K}} \mathbf{s}_\kappa (\mathbf{x}^{\iota_1^\kappa} - \mathbf{x}^{\iota_2^\kappa}) \boldsymbol{\alpha}_\kappa^s + \mathbf{B}^\top \boldsymbol{\beta}^s = \mathbf{x}^s \\
& \boldsymbol{\alpha}^s + \mu^s \mathbf{e} \leq \mathbf{0}
\end{aligned} \right\} \forall \mathbf{s} \in \mathcal{S}'. \tag{CCG_u^{\text{main}}(\mathcal{S}')}
\end{aligned}$$

This problem only involves a subset of the decision variables and constraints of Problem (1). Given variables $(\tau, \boldsymbol{\nu})$ feasible in the main problem, we define the \mathbf{s} th subproblem, $\mathbf{s} \in \tilde{\mathcal{S}}^K$, through

$$\begin{aligned}
& \text{maximize} && 0 \\
& \text{subject to} && \boldsymbol{\alpha} \in \mathbb{R}_+^K, \boldsymbol{\beta} \in \mathbb{R}_+^M, \mu \in \mathbb{R}_-, \mathbf{x} \in \mathcal{R} \\
& && \tau \leq \mathbf{b}^\top \boldsymbol{\beta} + \Gamma \mu \\
& && \sum_{\kappa \in \mathcal{K}} \mathbf{s}_\kappa (\mathbf{x}^{\iota_1^\kappa} - \mathbf{x}^{\iota_2^\kappa}) \boldsymbol{\alpha}_\kappa + \mathbf{B}^\top \boldsymbol{\beta} = \mathbf{x} \\
& && \boldsymbol{\alpha} + \mu \mathbf{e} \leq \mathbf{0}. \tag{CCG_u^{\text{sub},s}(\tau, \boldsymbol{\nu})}
\end{aligned}$$

To identify indices of subproblems $(\text{CCG}_u^{\text{sub},s}(\tau, \boldsymbol{\nu}))$ that, given a solution $(\tau, \boldsymbol{\nu})$ to the relaxed main problem, are infeasible, we solve a *single* feasibility MBLP defined through

$$\begin{aligned}
& \text{minimize} && \theta \\
& \text{subject to} && \theta \in \mathbb{R}, \mathbf{s} \in \tilde{\mathcal{S}}^K, \mathbf{u}^x \in \mathcal{U}^0, \boldsymbol{\epsilon}^x \in \mathcal{E}_\Gamma \quad \forall \mathbf{x} \in \mathcal{R} \\
& && \theta \geq (\mathbf{u}^x)^\top \mathbf{x} \quad \forall \mathbf{x} \in \mathcal{R} \\
& && \left. \begin{aligned}
& (\mathbf{u}^x)^\top (\mathbf{x}^{\iota_1^\kappa} - \mathbf{x}^{\iota_2^\kappa}) + \boldsymbol{\epsilon}_\kappa^x \geq M(\mathbf{s}_\kappa - 1) \\
& (\mathbf{u}^x)^\top (\mathbf{x}^{\iota_1^\kappa} - \mathbf{x}^{\iota_2^\kappa}) - \boldsymbol{\epsilon}_\kappa^x \leq M(\mathbf{s}_\kappa + 1)
\end{aligned} \right\} \forall \kappa \in \mathcal{K}, \mathbf{x} \in \mathcal{R}. \tag{CCG_u^{\text{feas}}(\boldsymbol{\nu})}
\end{aligned}$$

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

Proposition 1. *Let $\boldsymbol{\nu}$ be feasible in the relaxed master problem $(\text{CCG}_u^{\text{main}}(\mathcal{S}'))$. Then, $\boldsymbol{\nu}$ is feasible in Problem $(\widetilde{\mathcal{QP}}_u^K)$ and the objective value of $\boldsymbol{\nu}$ in Problem $(\widetilde{\mathcal{QP}}_u^K)$ is given by the optimal objective value of Problem $(\text{CCG}_u^{\text{feas}}(\boldsymbol{\nu}))$. If \mathcal{R} has fixed finite cardinality, then Problem $(\text{CCG}_u^{\text{main}}(\mathcal{S}'))$ can be solved as an MBLP using off-the-shelf solvers.*

Proposition 1 implies that, for any $\boldsymbol{\nu}$ feasible in the relaxed master problem $(\text{CCG}_u^{\text{main}}(\mathcal{S}'))$, the optimal value of $(\text{CCG}_u^{\text{feas}}(\boldsymbol{\nu}))$ yields a lower bound to the optimal value of Problem $(\widetilde{\mathcal{QP}}_u^K)$. At the same time, it is evident that for any index set $\mathcal{S}' \subseteq \tilde{\mathcal{S}}^K$, the optimal value of Problem $(\text{CCG}_u^{\text{main}}(\mathcal{S}'))$ yields an upper bound to the optimal objective value of Problem $(\widetilde{\mathcal{QP}}_u^K)$. The lemma below is key to identify indices of subproblems $\mathbf{s} \in \tilde{\mathcal{S}}^K$ that are infeasible.

Algorithm 1: Column-and-Constraint Generation procedure for solving Problem $(\widetilde{\mathcal{QP}}_u^K)$.

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 $\boldsymbol{\nu}^*$ from \mathcal{C}^K , near optimal in Problem $(\widetilde{\mathcal{QP}}_u^K)$ with associated objective θ ;

Initialization: $\boldsymbol{\nu}^* \leftarrow \emptyset$; 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 $(\text{CCG}_u^{\text{main}}(\mathcal{S}'))$. If it is solvable, let $(\tau, \boldsymbol{\nu}, \{\boldsymbol{\alpha}^s, \boldsymbol{\beta}^s, \boldsymbol{x}^s\}_{s \in \mathcal{S}'})$ be an optimal solution. If it is unbounded, set $\tau = \infty$ and let $\boldsymbol{\nu} \in \mathcal{C}^K$ be such that $(\text{CCG}_u^{\text{main}}(\mathcal{S}'))$ is unbounded when $\boldsymbol{\nu}$ is fixed to that value;

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

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

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

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

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

end

end

Set $\boldsymbol{\nu}^* \leftarrow \boldsymbol{\nu}$;

Result: Collection of queries $\boldsymbol{\nu}^*$ near-optimal in $(\widetilde{\mathcal{QP}}_u^K)$ with objective value θ .

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

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

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

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

Propositions 1 and Lemma 3 culminate in Algorithm 1 whose convergence is guaranteed by the following theorem.

Theorem 3. *Algorithm 1 terminates in a final number of steps with a feasible solution to Problem $(\widetilde{\mathcal{QP}}_u^K)$. The objective value attained by this solution is within δ of the optimal objective value of the problem.*

4.1.4. Speed-up Strategies We propose two additional speed-up strategies.

Symmetry Breaking. In Problem (\mathcal{QP}_u^K) , every permutation of the queries ι will yield another solution with the same objective function value. Correspondingly, Problem (2) also presents several symmetric solutions: we can permute the indices κ in the pairs $\{(\mathbf{v}^\kappa, \mathbf{w}^\kappa)\}_{\kappa \in \mathcal{K}}$ to build solutions with the same objective. To speed-up solution time, we propose to augment our MBLP formulation 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 constraints, note that each comparison is uniquely defined by the vector

$$\mathbf{y}^\kappa = \mathbf{v}^\kappa + \mathbf{w}^\kappa \quad (3)$$

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

$$\mathbf{z}^{\kappa\kappa'} \leq \mathbf{y}^\kappa + \mathbf{y}^{\kappa'}, \quad \mathbf{z}^{\kappa\kappa'} \leq 2 - \mathbf{y}^\kappa - \mathbf{y}^{\kappa'}, \quad \mathbf{z}^{\kappa\kappa'} \geq \mathbf{y}^\kappa - \mathbf{y}^{\kappa'}, \quad \mathbf{z}^{\kappa\kappa'} \geq \mathbf{y}^{\kappa'} - \mathbf{y}^\kappa \quad \forall \kappa, \kappa' \in \mathcal{K} : \kappa < \kappa'. \quad (4)$$

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

$$\mathbf{y}_i^{\kappa'} \geq \mathbf{y}_i^\kappa - \sum_{i' \in \mathcal{I} : i' < i} \mathbf{z}_{i'}^{\kappa\kappa'} \quad \forall i \in \mathcal{I}, \kappa, \kappa' \in \mathcal{K} : \kappa < \kappa'. \quad (5)$$

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

$$\sum_{i \in \mathcal{I}} \mathbf{z}_i^{\kappa, \kappa'} \geq 1 \quad \forall \kappa, \kappa' \in \mathcal{K} : \kappa < \kappa', \quad (6)$$

which eliminate solutions that involve asking the same queries twice.

As we will see in Section 4.3, the constraints above translate to significant speed-ups.

Warm-Starts. Solutions obtained by solving problems with smaller numbers of queries can be used to warm start problems with more queries. In this section, we describe how to build a feasible warm start for Problem (2) with $K + 1$ queries from a feasible solution to Problem (2) with K queries. Recall that a warm start need only fix solutions for the binary decision variables.

Let $(\tau, \mathbf{v}, \mathbf{w}, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{x}, \bar{\mathbf{v}}, \bar{\mathbf{w}})$ be feasible in Problem (2) with K queries. To generate a feasible solution to Problem (2) with $K + 1$ queries, we proceed as follows. First, we generate a new query ι^{K+1} randomly from the set

$$\mathcal{C} \setminus \left\{ \iota \in \mathcal{C} : \exists \kappa \in \mathcal{K} : \iota_1^\kappa = \sum_{i \in \mathcal{I}} i \cdot \mathbb{I}(\mathbf{v}_i^\kappa = 1) \text{ and } \iota_2^\kappa = \sum_{i \in \mathcal{I}} i \cdot \mathbb{I}(\mathbf{w}_i^\kappa = 1) \right\}$$

and define $\tilde{\mathbf{v}}$ and $\tilde{\mathbf{w}}$ through $\tilde{\mathbf{v}}^\kappa := \mathbf{v}^\kappa$, $\tilde{\mathbf{w}}^\kappa := \mathbf{w}^\kappa$, $\kappa \in \mathcal{K}$, $\tilde{\mathbf{v}}_i^{K+1} := \mathbb{I}(\boldsymbol{\iota}_1^{K+1} = i)$, and $\tilde{\mathbf{w}}_i^{K+1} := \mathbb{I}(\boldsymbol{\iota}_2^{K+1} = i)$ for each $i \in \mathcal{I}$. Subsequently, for each $\mathbf{s} \in \tilde{\mathcal{S}}^{K+1}$, we define $\tilde{\mathbf{x}}^{\mathbf{s}} := \mathbf{x}^{(\mathbf{s}_1, \dots, \mathbf{s}_K)}$. Then, $(\tilde{\mathbf{v}}, \tilde{\mathbf{w}}, \tilde{\mathbf{x}})$ constitutes a warm start for Problem (2) 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 (2) with $K + 1$ queries at least as high as that attained by $(\boldsymbol{\tau}, \mathbf{v}, \mathbf{w}, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{x}, \bar{\mathbf{v}}, \bar{\mathbf{w}})$ in Problem (2) 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 Section EC.2. In our numerical experiments, we will use a variant of this warm-start procedure which fixes the queries optimal in the problem with K queries in the problem with $K + 1$ queries to speed-up its solution. This approach is detailed in Section EC.3.

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

In Section 4.1, we assumed that all queries are chosen at once before the answers to any of the queries are observed. 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.

After making simplifications akin to those in Lemma 1 but adapted to this multi-stage setting, we can write the online active preference elicitation problem as

$$\max_{\boldsymbol{\iota}^1 \in \mathcal{C}} \min_{\mathbf{s}_1 \in \mathcal{S}} \max_{\boldsymbol{\iota}^2 \in \mathcal{C}} \min_{\mathbf{s}_2 \in \mathcal{S}} \cdots \max_{\boldsymbol{\iota}^K \in \mathcal{C}} \min_{\mathbf{s}_K \in \mathcal{S}} \max_{\mathbf{x} \in \mathcal{R}} \min_{\mathbf{u} \in \tilde{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \mathbf{s})} \mathbf{u}^\top \mathbf{x}. \quad (\mathcal{QP}_{\text{on}, \mathbf{u}}^K)$$

Indeed, for reasons similar to those in the proof of Lemma 1, we can allow responses \mathbf{s}^κ to be inconsistent with any $\mathbf{u} \in \tilde{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \mathbf{s})$ since this will not affect the optimal value nor the sets of optimal solutions of the query selection problems. Moreover, we can ignore indifferent responses without loss. In spite of these simplifications, Problem $(\mathcal{QP}_{\text{on}, \mathbf{u}}^K)$ is still significantly harder to solve than its offline counterpart (\mathcal{QP}_u^K) (or even $(\widetilde{\mathcal{QP}}_u^K)$).

We propose to leverage the tools we devised in Section 4.1 to obtain a conservative solution to the online problem. Specifically, we choose, at each step $k \in \mathcal{K}$, a query $\boldsymbol{\iota}^k$ as if it were the last query to be asked, i.e., in a one-step look-ahead fashion by solving problem (\mathcal{QP}_u^1) after updating its uncertainty set with the answers $\{\mathbf{s}_\kappa\}_{\kappa=1}^{k-1}$ obtained to the queries $\{\boldsymbol{\iota}^\kappa\}_{\kappa=1}^{k-1}$ actually asked. Concretely, we compute the k th query by solving

$$\text{maximize}_{\boldsymbol{\iota}^k \in \mathcal{C}} \min_{\mathbf{s}_k \in \tilde{\mathcal{S}}} \max_{\mathbf{x} \in \mathcal{R}} \min_{\mathbf{u} \in \tilde{\mathcal{U}}_{\Gamma(k)}(\boldsymbol{\iota}^{[k]}, \mathbf{s}_{[k]})} \mathbf{u}^\top \mathbf{x},$$

where, $\boldsymbol{\iota}^{[k]} := \{\boldsymbol{\iota}^\kappa\}_{\kappa=1}^k$, $\boldsymbol{s}_{[k]} := \{\boldsymbol{s}_\kappa\}_{\kappa=1}^k$, and with a slight abuse of notation

$$\tilde{\mathcal{U}}_{\Gamma(k)}(\boldsymbol{\iota}^{[k]}, \boldsymbol{s}_{[k]}) := \left\{ \begin{array}{l} \mathbf{u} \in \mathcal{U}^0 : \exists \boldsymbol{\epsilon} \in \mathcal{E}_{\Gamma(k)} \text{ such that} \\ \mathbf{u}^\top (\mathbf{x}^{\boldsymbol{\iota}_1^\kappa} - \mathbf{x}^{\boldsymbol{\iota}_2^\kappa}) \geq -\boldsymbol{\epsilon}_\kappa \quad \forall \kappa \in \{1, \dots, k\} : \boldsymbol{s}_\kappa = 1 \\ \mathbf{u}^\top (\mathbf{x}^{\boldsymbol{\iota}_1^\kappa} - \mathbf{x}^{\boldsymbol{\iota}_2^\kappa}) \leq \boldsymbol{\epsilon}_\kappa \quad \forall \kappa \in \{1, \dots, k\} : \tilde{\boldsymbol{s}}_\kappa = -1 \end{array} \right\}.$$

Note in particular that in the above definition, $\boldsymbol{\iota}^{[k-1]}$ and $\boldsymbol{s}_{[k-1]}$ are data. We propose to choose $\Gamma(k)$ in a fashion that exactly parallels the choice of Γ in the multi-stage problem.

Example 5 (Choice of $\Gamma(k)$). *If Γ is chosen as in Example 4, then a natural choice is $\Gamma(k) = 2\sigma\sqrt{k} \operatorname{erf}^{-1}(2p - 1)$ which ensures that $\boldsymbol{\epsilon} \in \mathcal{E}_{\Gamma(k)} \subseteq \mathbb{R}^k$ with probability p .*

4.3. Performance

We perform a wide range of experiments on randomly generated datasets of various sizes (I , J) and varying degree of inconsistencies (Γ) to evaluate the performance of our approach (MMU). To speed-up computation and to allow the solution of a large number of problems, unless stated otherwise, we employ the CCG algorithm from Section 4.1.3 which we augment with the symmetry breaking constraints from Section 4.1.4, and a greedy-heuristic that builds upon the warm-start approach, see Section EC.2. To make it easier to evaluate the performance of any given method, we normalize utilities so that they lie in the $[0, 1]$ range, where 0 corresponds to the worst-case utility if a recommendation is made before any questions are asked and 1 is the worst-case utility if the recommendation is made with full-knowledge of the utility. We provide here a summary of our results. Details of our experiments can be found in Section EC.5.

In the offline setting, we compare our approach to random elicitation (RAND) where the performance of a random query is evaluated by plugging it into Problem (\mathcal{QP}_r^K) (and using the methods in this paper to evaluate the objective of the resulting problem). The results are summarized in Table 1. From the table, it is apparent that computing an optimal set of queries using MMU is consistently preferable to selecting queries at random, resulting in an improvement in worst-case utility of over 63% on average across datasets and values of K . Moreover, we observe that while computing the worst-case utility of a set of queries (which is only possible thanks to our approach) is very fast (under 3 seconds on average), evaluating the quality of all sets of queries of cardinality K is computationally prohibitive (e.g., 2.3E+09 seconds for $K = 3$ on average). This motivates the use of our MMU approach that can compute optimal queries in practical times for the application (under 3 hours for realistic problem sizes). From the table it can also be noted that much of the useful information is gathered in the first 5 queries (increase in worst-case utility from 0.62 to 0.72). For completeness, we also conduct a computational study to evaluate the performance of

the CCG algorithm and symmetry breaking constraints. The results are summarized in Table 2. The table shows that relative to solving MBLP (2) directly, augmenting it with symmetry breaking constraints and using the CCG algorithm can yield significant speed-ups (up to 135 times on average). The table also shows that the heuristic used to speed-up computations is near-optimal, with an average optimality gap of under 2.1%.

Table 1 Performance of our offline elicitation procedure relative to random elicitation. All numbers correspond to averages across randomly generated datasets of various sizes. The numbers in the runtime columns correspond to the time it takes to compute an optimal query and to evaluate the performance of a random query (average over 50 sets of queries) in the case of **MMU** and **RAND**, respectively. The numbers in parenthesis in the last column correspond to the expected time it would take to evaluate the performance of all possible sets of K random queries.

K	Worst-Case Utility (Normalized)		Runtime (s)	
	MMU	RAND	MMU	RAND
1	0.62	0.25	35	1 (1.3E+03)
2	0.62	0.32	1668	1 (2.1E+06)
3	0.66	0.37	1854	1 (2.3E+09)
4	0.69	0.41	2346	1 (1.9E+12)
5	0.72	0.44	2936	1 (1.3E+15)
6	0.75	0.46	3370	2 (8.1E+17)
7	0.75	0.49	4435	2 (4.8E+20)
8	0.77	0.52	5759	2 (2.5E+23)
9	0.78	0.54	6637	2 (1.1E+26)
10	0.78	0.56	8121	3 (6.0E+28)
Avg.	0.71	0.44	3716	1.66 (6.0e+27)

Table 2 Performance of our speed-up methods and of the employed heuristic relative to solving the MBLP (2) directly. All numbers correspond to averages over randomly generated datasets.

Avg. speed-up of symmetry breaking	×65
Avg. speed-up of CCG+symmetry breaking	×135
Avg. optimality gap of heuristic	2.03%

In the online setting, we compare to random elicitation, to the polyhedral method of Toubia et al. (2004) (**POLY**), to the probabilistic polyhedral method of Toubia et al. (2007) (**PROB**), to the robust approach of Bertsimas and O’Hair (2013) (**ROB**), and to the ellipsoidal method of Sauré and Vielma (2019) (**ELL**). To evaluate the performance of any given method, we conduct simulations, wherein we draw 50 agents at random and simulate their responses (with and without inconsistencies). We record statistics on the guaranteed utility and rank of the recommended item which are summarized on Table 3. We omit solver times since all methods identify a query in under a second. From the table, it can be seen that our method consistently outperforms all other approaches in terms of both worst-case and average guaranteed utility (by over ×2.4 and ×1.5 on average, respectively). It

also results in improvements in worst-case and true rank of the recommendation (by over 1.8 and 1.2 spots on average, respectively). The table also shows that with MMU, most of the information is gained within the first 4-5 queries and that recommendations with almost optimal worst-case utility (0.91 on average) can be made after just 10 queries.

Table 3 Performance of our online elicitation procedure relative to state of the art methods from the literature over 50 randomly generated agents. All numbers correspond to averages across randomly generated datasets.

K	Worst-Case Guaranteed Utility (Normalized)						Worst-Case True Rank of Recommendation					
	MMU	RAND	POLY	PROB	ROB	ELL	MMU	RAND	POLY	PROB	ROB	ELL
1	0.62	0.05	0.35	0.35	0.12	0.58	27.71	35.57	32.86	32.86	37.86	26.71
2	0.74	0.15	0.16	0.19	0.20	0.26	24.43	34.71	33.43	32.14	36.71	33.00
3	0.81	0.22	0.20	0.17	0.26	0.20	21.57	32.14	32.00	33.57	35.29	32.57
4	0.84	0.27	0.18	0.19	0.30	0.20	19.71	30.86	31.71	30.00	35.00	29.71
5	0.87	0.36	0.15	0.16	0.34	0.26	17.57	29.00	28.57	29.43	33.57	29.00
6	0.89	0.38	0.18	0.19	0.36	0.29	16.43	29.29	28.14	29.43	30.71	26.71
7	0.89	0.42	0.16	0.20	0.37	0.31	16.57	27.43	29.14	27.43	29.43	26.43
8	0.90	0.50	0.16	0.18	0.36	0.37	14.43	27.29	28.86	30.29	29.00	25.86
9	0.91	0.52	0.27	0.29	0.37	0.42	13.71	25.29	27.43	27.14	28.43	22.43
10	0.91	0.54	0.25	0.28	0.40	0.50	13.86	25.43	28.43	27.29	29.86	19.43
Avg.	0.84	0.34	0.21	0.22	0.31	0.34	18.60	29.70	30.06	29.96	32.59	27.19

K	Average Guaranteed Utility (Normalized)						Average True Rank of Recommendation					
	MMU	RAND	POLY	PROB	ROB	ELL	MMU	RAND	POLY	PROB	ROB	ELL
1	0.62	0.23	0.36	0.36	0.17	0.59	11.59	13.72	14.11	14.13	14.43	11.35
2	0.75	0.34	0.38	0.40	0.29	0.49	9.19	12.65	12.83	11.44	13.83	11.26
3	0.82	0.40	0.40	0.41	0.36	0.43	7.17	11.88	11.25	10.58	11.87	10.86
4	0.86	0.47	0.43	0.47	0.41	0.44	6.31	11.04	11.46	8.49	10.90	8.86
5	0.89	0.52	0.47	0.50	0.45	0.52	5.65	10.26	7.96	8.07	10.05	7.53
6	0.91	0.56	0.53	0.55	0.50	0.55	5.23	9.20	7.81	7.21	9.36	6.58
7	0.92	0.61	0.54	0.56	0.53	0.59	5.18	8.43	7.17	6.76	8.68	6.08
8	0.93	0.65	0.58	0.61	0.54	0.64	4.77	7.56	6.99	6.47	8.37	5.08
9	0.93	0.68	0.62	0.63	0.56	0.70	4.39	7.18	6.28	6.26	7.85	4.62
10	0.94	0.70	0.64	0.66	0.59	0.73	4.32	7.05	6.71	5.73	7.82	4.00
Avg.	0.86	0.54	0.49	0.51	0.44	0.57	6.38	9.72	9.26	8.51	10.32	7.62

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

We now study the offline and online active preference elicitation problems under the min-max regret decision criterion and evaluate their performance on synthetic data.

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

5.1.1. Problem Formulation, Complexity Analysis, & MBLP Reformulation. The offline active preference elicitation problem under the min-max regret decision criterion is a variant of its max-min utility counterpart where the objective of the recommender system when selecting queries and making recommendations is to minimize worst-case regret (rather than maximize worst-case

utility). Mathematically, it is expressible as the following two-stage robust optimization problem with decision-dependent information discovery

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

Problem (\mathcal{QP}_r^K) is generally \mathcal{NP} -hard when the recommendation set is discrete as summarized in the following theorem.

Theorem 4. *Evaluating the objective function of Problem (\mathcal{QP}_r^K) is \mathcal{NP} -hard even if the recommendation set \mathcal{R} consists of only two elements.*

We note that in contrast to the max-min utility setting, since the objective function of the min-max regret recommendation problem is not concave in \boldsymbol{u} , the minimax theorem that could be used to simplify the max-min utility problem when \mathcal{R} is convex (see Theorem 1) does not apply here. This result motivates us to reformulate Problem (\mathcal{QP}_r^K) as an MBLP as shown in the following theorem. Throughout the remainder of this section, we assume that \mathcal{R} has fixed finite cardinality.

Theorem 5. *Problem (\mathcal{QP}_r^K) is equivalent to the following mixed-binary linear program*

$$\begin{aligned} & \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}^s \in \mathcal{R} \quad \forall \boldsymbol{s} \in \tilde{\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 \tilde{\mathcal{S}}^K, \boldsymbol{x}' \in \mathcal{R} \\ & && \bar{\boldsymbol{v}}^{(\boldsymbol{x}', \boldsymbol{s})\kappa}, \bar{\boldsymbol{w}}^{(\boldsymbol{x}', \boldsymbol{s})\kappa} \in \mathbb{R}_-^I \quad \forall \boldsymbol{s} \in \tilde{\mathcal{S}}^K, \boldsymbol{x}' \in \mathcal{R}, \kappa \in \mathcal{K} \\ & && \left. \begin{aligned} \tau &\geq \boldsymbol{b}^\top \boldsymbol{\beta}^{(\boldsymbol{x}', \boldsymbol{s})} + \Gamma \mu^{(\boldsymbol{x}', \boldsymbol{s})} \\ \sum_{i \in \mathcal{I}} \boldsymbol{x}^i \sum_{\kappa \in \mathcal{K}} s_\kappa \left(\bar{\boldsymbol{v}}_i^{(\boldsymbol{x}', \boldsymbol{s})\kappa} - \bar{\boldsymbol{w}}_i^{(\boldsymbol{x}', \boldsymbol{s})\kappa} \right) + \boldsymbol{B}^\top \boldsymbol{\beta}^{(\boldsymbol{x}', \boldsymbol{s})} &= \boldsymbol{x}' - \boldsymbol{x}^s \\ \boldsymbol{\alpha}^{(\boldsymbol{x}', \boldsymbol{s})} + \mu^{(\boldsymbol{x}', \boldsymbol{s})} \boldsymbol{e} &\geq \mathbf{0} \end{aligned} \right\} \quad \left. \begin{aligned} \forall \boldsymbol{s} \in \tilde{\mathcal{S}}^K, \\ \boldsymbol{x}' \in \mathcal{R} \end{aligned} \right\} \quad (7) \\ & && \left. \begin{aligned} \boldsymbol{e}^\top \boldsymbol{v}^\kappa = 1, \boldsymbol{e}^\top \boldsymbol{w}^\kappa = 1, 1 - \boldsymbol{w}_i^\kappa &\geq \sum_{i': i' \geq i} \boldsymbol{v}_{i'}^\kappa \quad \forall i \in \mathcal{I}, \forall \kappa \in \mathcal{K} \\ \bar{\boldsymbol{v}}^{(\boldsymbol{x}', \boldsymbol{s})\kappa} &\geq -M \boldsymbol{v}^\kappa, \bar{\boldsymbol{v}}^{(\boldsymbol{x}', \boldsymbol{s})\kappa} &\geq \boldsymbol{\alpha}_\kappa^{(\boldsymbol{x}', \boldsymbol{s})} \boldsymbol{e} \\ \bar{\boldsymbol{v}}^{(\boldsymbol{x}', \boldsymbol{s})\kappa} &\leq \boldsymbol{\alpha}_\kappa^{(\boldsymbol{x}', \boldsymbol{s})} \boldsymbol{e} + M(\boldsymbol{e} - \boldsymbol{v}^\kappa) \\ \bar{\boldsymbol{w}}^{(\boldsymbol{x}', \boldsymbol{s})\kappa} &\geq -M \boldsymbol{w}^\kappa, \bar{\boldsymbol{w}}^{(\boldsymbol{x}', \boldsymbol{s})\kappa} &\geq \boldsymbol{\alpha}_\kappa^{(\boldsymbol{x}', \boldsymbol{s})} \boldsymbol{e} \\ \bar{\boldsymbol{w}}^{(\boldsymbol{x}', \boldsymbol{s})\kappa} &\leq \boldsymbol{\alpha}_\kappa^{(\boldsymbol{x}', \boldsymbol{s})} \boldsymbol{e} + M(\boldsymbol{e} - \boldsymbol{w}^\kappa) \end{aligned} \right\} \quad \left. \begin{aligned} \forall \boldsymbol{s} \in \tilde{\mathcal{S}}^K, \boldsymbol{x}' \in \mathcal{R}, \\ \kappa \in \mathcal{K}, \end{aligned} \right\} \end{aligned}$$

where M is a “big- M ” constant. In particular, given an optimal solution $(\tau, \boldsymbol{v}, \boldsymbol{w}, \boldsymbol{\alpha}, \boldsymbol{\beta}, \bar{\boldsymbol{v}}, \bar{\boldsymbol{w}})$ to Problem (7), an optimal set of queries for Problem (\mathcal{QP}_r^K) is given by

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

5.1.2. Column-and-Constraint Generation For fixed K , Problem (7) is polynomial in the size of the input. Yet, it presents a number of decision variables and constraints that are exponential in K . We thus propose a column-and-constraint generation approach, similar to the one proposed in Section 4.1.3 for the max-min utility case, that enables us to scale to practical values of K . In the interest of space, we defer the details to the electronic companion, see Section EC.4.

Naturally, the speed-up strategies from Section 4.1.4 directly apply to the min-max regret setting.

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

Similar to the max-min utility case, we can formulate the online version of the active preference elicitation problem with the min-max regret decision criterion as a min-max-...-min-max-max problem with decision-dependent uncertainty sets. We propose to solve this problem approximately in a one-step look-ahead fashion in a way that exactly parallels the approach from Section 4.2.

5.3. Performance

We perform a wide set of experiments on randomly generated data to evaluate the performance of our min-max regret elicitation procedure (MMR). Our results are summarized in Table 4. Since these are qualitatively very similar to those reported in the max-min utility case, see Section 4.3, we only report averages across datasets and values of K and refer the reader to Section EC.5 for more information. We note that since the regret problem is harder than the utility elicitation problem, the instances considered here were smaller than those in Section 4.3 to be able to report results on a variety of problems. From the table, we observe that our proposed approach outperforms existing methods on all metrics on average. Most notably, our offline method reduces the worst-case regret by a factor of 2 on average. Although the reported numbers are averages, a similar behavior is exhibited by individual realizations.

Table 4 Performance of our min-max regret offline (top two rows) and online (bottom four rows) elicitation procedures relative to state of the art methods from the literature. All numbers correspond to averages across K and randomly generated datasets; for the online elicitation case, we use 50 randomly generated agents.

	MMR	RAND	POLY	PROB	ROB	ELL
Worst-Case Regret (Normalized)	0.51	0.66	–	–	–	–
Runtime (s)	1540	2 (1E+20)	–	–	–	–
Worst-Case Guaranteed Regret (Normalized)	0.38	0.77	0.82	0.83	0.72	0.58
Worst-Case True Rank of Recommendation	4.37	7.74	8.04	8.33	7.69	5.57
Average Guaranteed Regret (Normalized)	0.31	0.54	0.59	0.62	0.53	0.41
Average True Rank of Recommendation	1.72	2.56	2.65	2.66	2.74	1.95

In our last set of experiments on synthetic data, we investigate whether there are benefits in employing the min-max regret solution relative the max-min utility solution – in particular, whether

it results in solutions that have lower regret than solutions to the problem that maximizes worst-case utility. The results are summarized on Table 5, where it can be seen that the worst-case regret queries (worst-case utility) have lower regret (higher utility) than their max-min utility (min-max regret) counterparts. This implies that although computing min-max regret queries comes at a higher computational cost, they may be worthwhile to compute if one is keen on having low regret.

Table 5 Comparison between max-min utility and min-max regret based queries on a randomly generated synthetic dataset. All approaches were given a one hour time limit. The numbers in the table correspond to averages across values of K .

Max-Min Utility Solution		Min-Max Regret Solution	
Normalized Worst-Case Utility	Normalized Worst-Case Regret	Normalized Worst-Case Utility	Normalized Worst-Case Regret
0.84	0.58	0.75	0.53

6. Designing Allocation Policies that Align with Human Value Judgements

In this section, we investigate a real-world application of our preference elicitation framework. In particular, we consider the problem of eliciting the preferences of policy-makers in charge of allocating scarce housing resources to people experiencing homelessness. We envision our methodology will enable homeless services organizations such as LAHSA to allocate resources in a way that aligns with the values of its policy-makers to adequately balance the interests of those they serve. We first describe the L.A. housing allocation system and motivate various attributes of policies that are relevant to be considered when evaluating them against one another. We then discuss how the dataset we have at our disposal can be used to estimate the performance of counterfactual housing allocation policies in terms of these metrics. We conclude by performing a numerical case study based on historical data that showcases the promise of the proposed approach.

6.1. The L.A. Housing Allocation System

In LA County, LAHSA coordinates housing and services for homeless families and individuals. It employs a *Coordinated Entry System (CES)* which facilitates the coordination and management of resources and services. The CES enables organizations within the same community to pool both their housing resources and clients. The County is divided into eight geographic areas designated as Service Planning Areas (SPAs). Each SPA is expected to have a balance of homeless services used to serve clients in its area and consists of three systems that operate mostly independently of one another: the Youth system that serves those under 24 years old, the Single Adult system, and the Family system. Each SPA System has a “matcher” in charge of working with LAHSA and with case workers that know the clients to match resources as they become available to waitlisted

clients. For simplicity, and to align with the data we have at our disposal, we focus here on the Single Adult system and note that the other systems operate in a similar fashion.

When an individual interacts with the housing allocation system for the first time, e.g., through street outreach or by staying at a shelter, they typically answer a survey, the Vulnerability Index – Service Prioritization Decision Assistance Tool (VI-SPDAT), which will help determine risk and prioritization when providing them with assistance (Orgcode 2014). The VI-SPDAT consists of 34 predominantly yes-or-no items intended to measure an individual’s level of vulnerability across four domains: their history of housing and homelessness, individual risk factors, socialization and daily functions, and wellness. The answers to the VI-SPDAT survey are also used to compute a *risk score* ranging from 0 to 17 with higher scores indicating greater vulnerability. This information is entered into the Homeless Management Information System (HMIS) together with client demographic information (e.g., age, race/ethnicity). At this point, they are waitlisted for housing in their SPA and matchers can see them in their system, in what is called the *community queue*.

Housing resources can broadly be categorized in two groups: *permanent supportive housing (PSH)*, which combines low-barrier affordable housing, health care, and other supportive services; and *rapid rehousing (RRH)*, which provides short-term rental assistance and services. Individuals who do not receive a housing resource will still receive services depending on their needs.

To help decide how to match people to resources, a housing allocation policy is in place which, roughly speaking, instructs matchers to direct clients towards more supportive resources the more vulnerable they are (according to the score). In LA, the allocation policy recommends that those who score 12 or higher receive PSH, those who score 8-11 receive either PSH or RRH, and those who score 7 or lower receive RRH. Different communities employ variants of this policy with different “cut-scores” and potentially having a bucket of (low) scores for which no resource is intended. A salient advantage of this policy is its *interpretability* which is very important when operating in high-stakes settings (Rudin 2019). Yet, this policy is not tied to outcomes in the sense that it has not been designed with e.g., the goal of being effective nor, in some sense, fair.

In practice, this recommended policy is not fully adhered to due to several reasons of which we discuss the main ones. First, it is not implementable as there are insufficient resources to match the demand in each score “bucket”. Second, the matcher and case worker, when discussing a particular case during *case conferencing*, may deem that an individual’s vulnerability score is not representative of their true vulnerability. Third, resources are matched as they arrive and would never “wait” for a more vulnerable individual. Finally, to date matchers are only responsible for matching PSH resources whereas RRH resources are matched by referral through a mechanism that is more ad hoc (this is in the process of changing).

Each time a resource is allocated to an individual, the match is reported in the HMIS together with all interactions that the individual had with the system. In particular, in addition to PSH/RRH matches, other exits from homelessness are recorded, such as if an individual self-resolved (independently obtained private housing) or went to live with family.

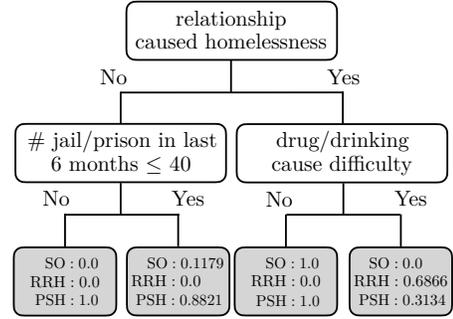
Being sensitive to the fact that it supports communities disproportionately affected by systemic racial, gender, and other injustices, LAHSA instigated a number of committees in charge of evaluating their system from an equity lens. Most notably, it assembled the Ad Hoc Committee on Women and Homelessness and the Ad Hoc Committee on Black People Experiencing Homelessness (LAHSA 2017, 2018). An important finding of the latter committee is that in the current system, although resources are allocated at equal rates among Black and White people, Black people return to homelessness at twice the rates of Whites. Similar findings have been reported in other communities (Petry et al. 2021, Hill et al. 2021, Hsu et al. 2021). Motivated by these findings and by the desire to make more effective use of its resources, LAHSA is interested in building a more evidence based system that will effectively allocate resources while better aligning with the values of policy-makers in serving its clients with a focus on equity.

6.2. Trade-Offs and Policy Attributes

A housing allocation policy is a function $\pi : \mathcal{X} \rightarrow \Delta^{|\mathcal{T}|}$ that maps an individual’s characteristics $x \in \mathcal{X}$ to their probability of getting each intervention in the set $\mathcal{T} := \{\text{PSH}, \text{RRH}, \text{SO}\}$, where SO stands for “Services Only”, i.e., no housing intervention. We focus our attention on policies that are *feasible*, in the sense that the expected number of clients we assign to PSH and RRH does not exceed the number of these resources that are available (according to this definition, the current recommended policy is not feasible). From the above discussion, the key characteristics that drive the choice of a policy are its interpretability, effectiveness, and fairness properties. We now formalize these mathematically.

Prescriptive Trees and Measures of Interpretability. The first step in designing a policy is agreeing on one (or several) policy classes to consider (Azizi et al. 2018). Motivated by the desire to have interpretable policies, we focus here on prescriptive trees (Kallus 2017, Bertsimas et al. 2019, Jo et al. 2021). These take the form of a binary tree. In each branching node of the tree, a binary test is performed on a specific feature. Two branches emanate from each branching node, with each branch representing the outcome of the test. If a datapoint passes (resp. fails) the test, it is directed to the right (resp. left) branch. At each leaf node, each treatment is assigned with a certain probability. Thus, each path from root to leaf represents a rule that assigns each treatment with the same probability to all individuals that reach that leaf. An example of a prescriptive tree can be seen in Figure 1. Prescriptive trees are among the most interpretable treatment rules and

Figure 1 Sample prescriptive tree to allocate housing resources based on answers to the VI-SPDAT survey questions. The numbers at each leaf correspond to the probability of getting each intervention for all individuals that reach that leaf.



are thus particularly attractive when designing social and public health interventions. We measure interpretability of a prescriptive tree using two attributes: the number of branching nodes and the number of features it uses (the greater these are, the less interpretable the policy).

Measure of Effectiveness. To measure effectiveness of a policy, we must use outcomes that policy-makers care about and that are also observable in the HMIS. According to the U.S. Department of Housing and Urban Development (HUD), an important outcome to mitigate is *chronic homelessness*. While this is not directly observable in the data, a proxy that can be observed is *housing (in)stability*. Using potential outcomes notations, see Hernán and Robins (2019), we let $Y(t) \in \{0, 1\}$ denote the outcome under treatment $t \in \mathcal{T}$ —thus $Y(t) = 1$ iff the individual will still be housed after one year of receiving treatment t . Note that individuals who do not receive a PSH nor RRH resource may still have a positive outcome if e.g., they self-resolved. Mathematically, we can express effectiveness as the expected number of people that are still housed under the given policy, i.e.,

$$Q(\pi) := \mathbb{E} \left[\sum_{t \in \mathcal{T}} \pi(t|X)Y(t) \right], \tag{8}$$

where the expectation is taken over the joint distribution of $X, Y(1), \dots, Y(T)$.

Measures of Fairness. To measure fairness, we focus on metrics that have been investigated by the ad hoc committees charged with evaluating equity in the system, see Section 6.1. These look at outcomes by group with respect to various protected characteristics. Letting \mathcal{P} denote the set of protected characteristics and \mathcal{X}_p the set of values that this covariate can take, we can express the conditional average outcome for each group as

$$Q_p(\pi, x_p) := \mathbb{E} \left[\sum_{t \in \mathcal{T}} \pi(t|X)Y(t) \middle| X_p = x_p \right] \text{ for all } p \in \mathcal{P} \text{ and } x_p \in \mathcal{X}_p, \tag{9}$$

where $X_p \subseteq X$ represents the value of protected characteristic p . An associated metric that we also consider an important attribute in the choice of policy is the smallest conditional average outcome for each protected characteristic, i.e., $\min_{x_p \in \mathcal{X}_p} Q_p(\pi, x_p)$ for all $p \in \mathcal{P}$. This is motivated by the notion of max-min fairness that seeks to provide maximum benefit to the least advantaged members of society (Rawls 1991). The last attribute related to fairness that we consider is the number of protected features we branch on (with smaller numbers being preferred).

6.3. Using our Preference Elicitation Framework

To use our framework in practice several candidate policies must be identified and their fairness, effectiveness, and interpretability attributes must either be calculated (e.g., number of branching nodes) or estimated from data (e.g., outcomes by race). We now describe in more detail the data we have and the process we use to estimate the various fairness and effectiveness metrics.

Raw Data Description. The raw dataset we use features HMIS data from 16 Continuums of Care that were sourced, accessed, anonymized, and provided to us by OrgCode Consulting, Inc. The dataset includes the demographic characteristics and VI-SPDAT responses for $I = 22,165$ unsheltered single adults assessed between February 2015 and April 2018. Demographic characteristics include individuals' self-reported age, gender, and race or ethnicity. We collect demographic information and VI-SPDAT responses in the feature vector $X_i \in \mathcal{X} \subseteq \mathbb{R}^F$, $i \in \mathcal{I} := \{1, \dots, I\}$. The data also includes information on the intervention $T_i \in \mathcal{T}$ received by individual $i \in \mathcal{I}$. Housing interventions include exits from homelessness to either PSH, RRH, family, or self-resolve. In our analysis, we group family and self-resolve interventions under the SO category. Finally, the data reports the observed outcome Y_i , which indicated if the client is still housed 365 days after their initial exit from homelessness, see Petry et al. (2021) for details..

Protected Characteristics and Groups. We measure fairness with respect to race/ethnicity, age, and gender. We group individuals into three racial/ethnic groups, namely White, Black, and Other, where the Other category represents 5.9% of the population and includes those self reported as Asian, Hispanic, Native American, and Pacific Islander. Similarly, we group individuals into four age groups, those aged [25,41], (41,48], (48,54], and (54,84], such that each group has approximately the same size. Finally, there were three self reported gender groups—male, female, and transgender. Thus, $\mathcal{P} := \{\text{race, gender, age}\}$ and e.g., $\mathcal{X}_{\text{race}} := \{\text{White, Black, Other}\}$.

Policy Attributes. This approach of grouping individuals gives us a total of $J = 17$ features for each policy: 1 related to efficiency, 2 to interpretability, and the remaining to fairness.

Potential Outcomes and Assumptions. The HMIS data is an observational dataset where housing interventions are not allocated at random (as in randomized clinical trials). We view the HMIS data in a potential outcomes framework. In particular, the joint distribution of X and $\{Y(t)\}_{t \in \mathcal{T}}$ is unknown and we can only observe, for each point i in the data the outcome $Y_i = Y_i(T_i)$ under the treatment received. Critically, we cannot control (and do not know) the historical treatment assignment policy and the outcomes $Y_i(t)$ for $t \neq T_i$ remain *unobserved*. We assume however, that the conditional probability of receiving a treatment depends only on the covariates X and that the probability of receiving any given treatment conditional on X is positive, i.e., $\mu(t, x) := \mathbb{P}(T = t | X = x) > 0$ almost surely for all t . These two assumptions, standard in causal inference, will

allow us to use the available data to estimate the performance of counterfactual housing allocation policies. While not possible to verify them from the data, these assumptions mostly hold based on what we know about the system (indeed, while some data from case conferencing that is not in HMIS may have been used to assign resources, this seems to be a rare occurrence).

Estimating the Performance of Counterfactual Policies. Evaluating several of the fairness and efficiency metrics such as (8) and (9) requires knowledge of the joint distribution of X and $\{Y(t)\}_{t \in \mathcal{T}}$. While this distribution is unknown, we can leverage the assumptions above to estimate them from the historical observations $\{(X_i, T_i, Y_i)\}_{i \in \mathcal{I}}$. In our work, we propose to use a doubly robust estimator, see Dudík et al. (2011). Specifically, we estimate $Q(\pi)$ and $Q_p(\pi, x_p)$ through

$$\hat{Q}(\pi) := \frac{1}{I} \sum_{i \in \mathcal{I}} \left(\hat{\nu}_{\pi(X_i)}(X_i) + (Y_i - \hat{\nu}_{T_i}(X_i)) \frac{\pi(T_i|X_i)}{\hat{\mu}(T_i, X_i)} \right) \quad (10)$$

and

$$\hat{Q}_p(\pi, x_p) := \frac{1}{|\{i \in \mathcal{I} : X_{i,p} = x_p\}|} \sum_{i \in \mathcal{I} : X_{i,p} = x_p} \left(\hat{\nu}_{\pi(X_i)}(X_i) + (Y_i - \hat{\nu}_{T_i}(X_i)) \frac{\pi(T_i|X_i)}{\hat{\mu}(T_i, X_i)} \right), \quad (11)$$

respectively. Here, $\hat{\mu}$ is an estimator of μ , which is obtained for instance using machine learning, by fitting a model to $\{(X_i, T_i)\}_{i \in \mathcal{I}}$, and $\hat{\nu}_t(x)$ is an estimator of $\mathbb{E}(Y|T = t, X = x)$ obtained by fitting a model on the subpopulation that was assigned treatment t . Provided at least one of $\hat{\mu}$ or $\hat{\nu}$ converges almost surely to μ or ν , respectively, these doubly robust estimators of $Q(\pi)$ and $\hat{Q}_p(\pi, x_p)$ will be asymptotically unbiased. All fairness and efficiency metrics of a policy π can then be estimated using the two estimators above.

6.4. Numerical Case Study

We now summarize our findings. We first discuss our models of $\hat{\nu}_t(x)$, $t \in \mathcal{T}$, and $\hat{\mu}(t, x)$ and then present the results of the offline and online elicitation procedures.

Since we use estimation procedures to evaluate the counterfactual performance of a policy and in particular its fairness properties, it is imperative that our learned models of $\hat{\nu}_t(x)$ and $\hat{\mu}(t, x)$ in (10) and (11) be fair. In particular, they should have the same true positive rates (TPR) and same false positive rates (FPR) across groups for each protected group and feature. This notion of fairness is known as *equalized odds*, see Hardt et al. (2016). We calculate the equalized odds fairness metric by taking the maximum, across all protected characteristics, of the greatest difference in both TPR and FPR between all groups. Thus, this is a number between 0 and 1 and we would like it to be as small as possible to indicate a fair classifier. We use the approach from Aghaei et al. (2020) to learn optimal fair classification trees—we denote this by **OFT**. This framework has the advantage that it performs close to the best unfair models (e.g., CART, random forests) in terms of accuracy when no fairness constraints are imposed, that it applies to multi-class classification problems, and to problems with multiple protected groups and multiple levels for each group.

Fair Learning of $\hat{\nu}_t(x)$. We used several models to learn outcomes $\hat{\nu}_t(x) = \mathbb{E}(Y|X = x, T = t)$ for $t \in \mathcal{T}$. Results of the best models are shown on Figure 2. As shown in the figure, in both training and testing the best CART model had equalized odds of over 75% showcasing the importance of imposing fairness directly in the model. In contrast, by using OFT, we can see that in all cases we could significantly decrease equalized odds of the learned models compared to CART at no cost to accuracy. For $t \in \{\text{SO}, \text{RRH}\}$, decreasing equalized odds beyond about 20% results in significant drop in accuracy. For our predictions, we chose the model associated with the lowest possible equalized odds that did not result in significant drop in accuracy. For the case of the SO model, rather than optimizing accuracy, we optimized *balanced* accuracy since this dataset was significantly imbalanced. From the figure, we can see that the learned models are interpretable and also make a lot of sense. For example, in the case of the SO intervention, only those recently stably housed with no more than one episode of homelessness are predicted to have a positive outcome.

Fair Learning of $\hat{\mu}(t, x)$. Learning the propensity score is a multi-class classification problem. We generalize the equalized odds fairness metric to multiple classes as the greatest difference in error rates of any type across groups. Our results on learning $\hat{\mu}(t, x)$ are summarized in Figure 3. From the figure, we see that the best CART model results in equalized odds of over 80%, whereas using OFT allows us to carefully tune the trade-off between fairness and accuracy and in particular to reduce equalized odds to under 38% at no cost to accuracy. Moreover, we see that the model obtained to fairly predict propensity score is very interpretable and natural: the more the indications of vulnerability an individual has, the more likely they are to get more supportive interventions.

Generating Feasible Counterfactual Prescriptive Tree-Based Policies. To generate feasible counterfactual prescriptive trees with different fairness, efficiency, and interpretability properties, we proceed as follows. For each policy we wish to design (25 in total), we sample a utility vector \mathbf{u} uniformly at random from the set $\mathcal{U} = \{\mathbf{u} \in \mathbb{R}_+^{17} : \mathbf{e}^\top \mathbf{u} = 1\}$. We also draw the depth d of the tree uniformly from the set $\{1, 2, 3\}$. Then, we use the method from Jo et al. (2021) to design an optimal and feasible tree based policy of depth d that optimizes the combination of the aforementioned metrics of fairness, efficiency, and interpretability with weight vector \mathbf{u} . Finally, we record the relevant attributes of the policy, resulting in a dataset with $I = 25$ policies and $J = 17$ attributes.

Results of Offline MMR Elicitation. We focus on the min-max regret case as this usually results in less conservative solutions. The results of offline MMR on this dataset of candidate prescriptive tree policies are shown on Figure 4. From the figure, we see that our approach consistently outperforms random elicitation across all numbers of queries in terms of worst-case regret. In particular, with just a single query, our approach is able to outperform the average performance of random elicitation with 10 queries. This is true whether responses given by the user are inconsistent or not. In addition, optimizing over queries is faster than evaluating the performance of all possible queries.

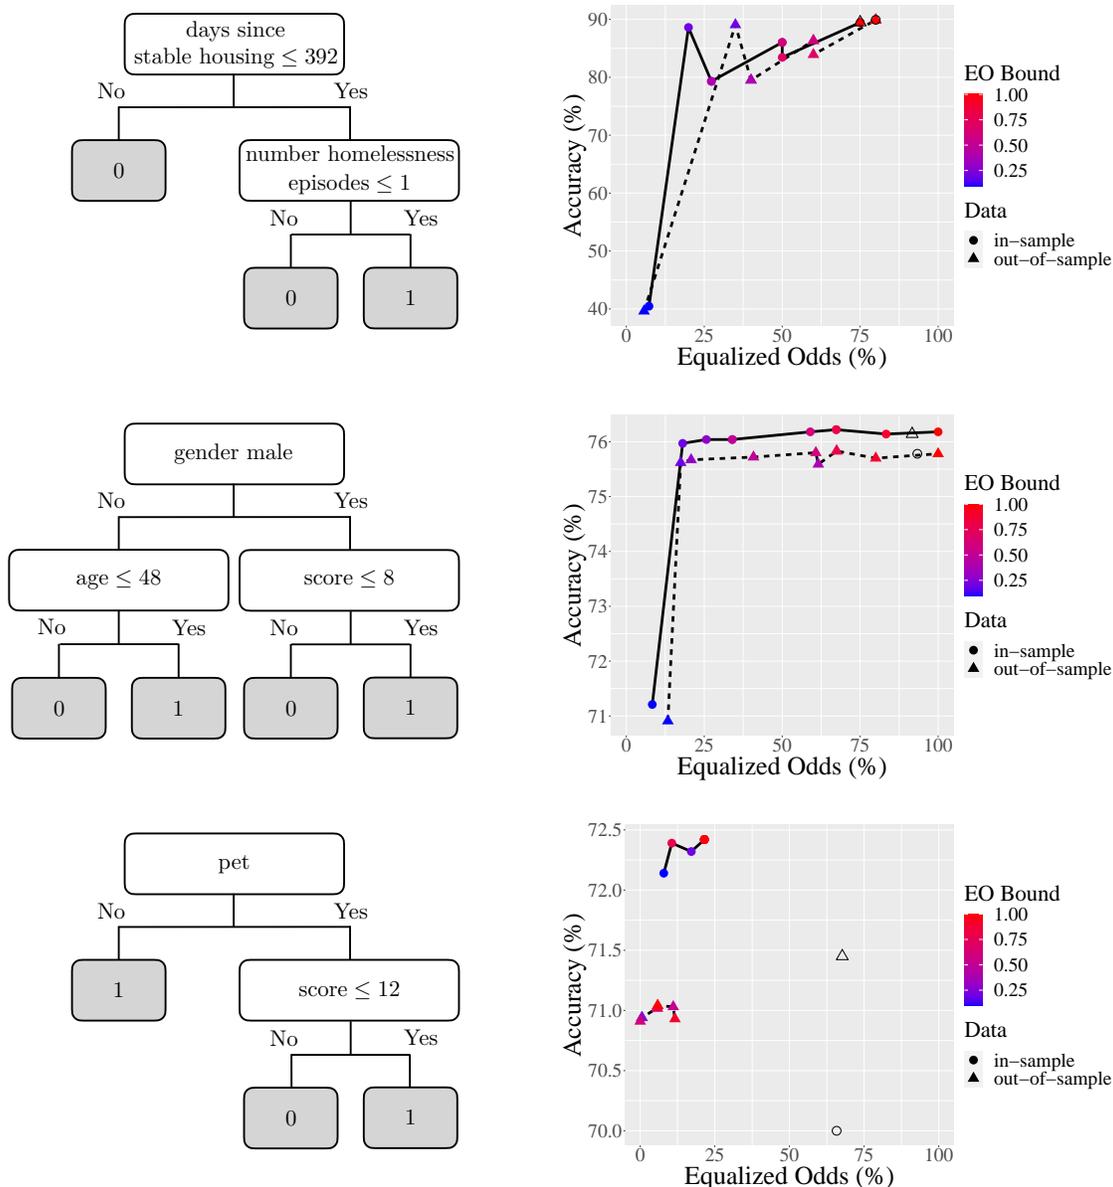


Figure 2 Optimal and fair decision trees for estimating $\nu_t(x) = \mathbb{E}(Y|X = x, T = t)$ for $t \in \{\text{SO}, \text{RRH}, \text{PSH}\}$ (top, middle, and bottom, respectively). The figures on the right show the fairness/accuracy trade-off both in-sample and out-of-sample for CART (empty circle and triangle) and for OFT (filled circles and triangles). The figures on the left show the optimal fair classification trees that we chose to use; these corresponded to the trees that were the most fair among those that resulted in minimal loss in accuracy.

Results of Online MMR Elicitation. Since in the online case it is not possible to compute the objective value of the problem, we evaluate our method through simulation by sampling 50 agent utilities from the uncertainty set and allowing potential inconsistencies in their responses. The results are summarized in Figure 5. From the figure, it can be seen that our approach consistently outperforms all other methods from the literature in terms of worst-case regret, worst-case true

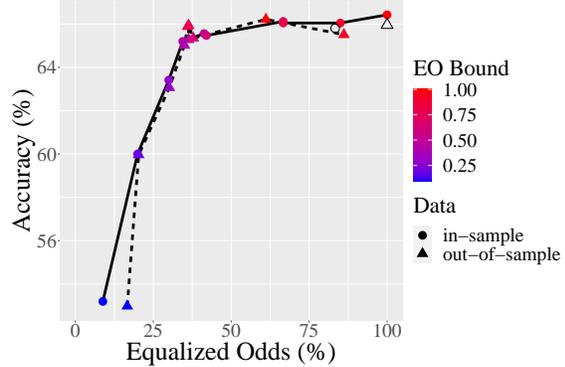
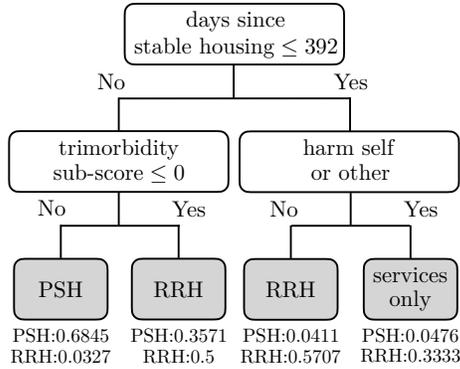
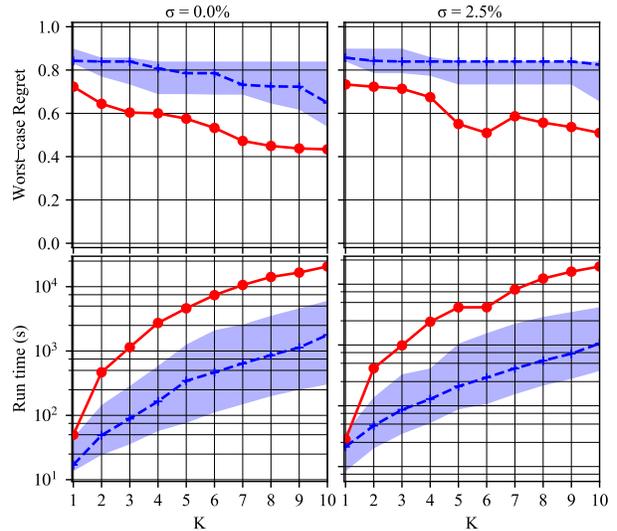


Figure 3 Optimal and fair decision-tree for predicting $\hat{\mu}(t, x)$. The meaning of the circles and triangles is the same as in Figure 2. The tree chosen to predict propensity scores is shown on the left.

Figure 4 Optimality (top) and scalability (bottom) results for the offline min-max regret preference elicitation problem (\mathcal{QP}_r^K) on the real data when $\sigma = 0$ (left) and $\sigma = 0.025$ (right). Approach MMR is shown with red dots. The median performance of **RAND** across 50 sets of K random queries is shown with the blue dashed line. In particular the run time for **RAND** is the time it takes to evaluate the performance of one query. The blue shaded region shows the range of performance of **RAND** across these 50 sets of queries.



regret of the recommendation, and worst-case true rank of the recommendation. In particular, after just 7 queries, our method is able to consistently recommend the best item if there are no inconsistencies and an item ranked 6 or higher if there are inconsistencies.

Robustness to Choice of Γ . A challenge when working with real data is that our model may be misspecified, e.g., if the standard deviation σ used to select queries is different than the true standard deviation σ^* of the shocks in the user utility, see Section 3. We perform an evaluation of our model on a range of such misspecifications. The results are summarized in Table 6. From the table, we see that our query selection methods are quite robust to our modeling assumptions. In particular, worst-case regret, worst-case true regret, and worst-case true rank only mildly deteriorate if the model is different than assumed.

Figure 5 Worst-case regret (top), worst-case true regret (middle), and worst-case rank (bottom) for the online min-max regret preference elicitation problem on real data when $\sigma = 0$ (left) and $\sigma = 0.025$ (right). The median performance of MMR (resp. RAND, POLY, PROB, ROB, ELL) across 50 random utility vectors \mathbf{u} and inconsistencies ϵ is shown with red dots (resp. blue plus, black triangles, black inverted triangles, green crosses, and purple squares).

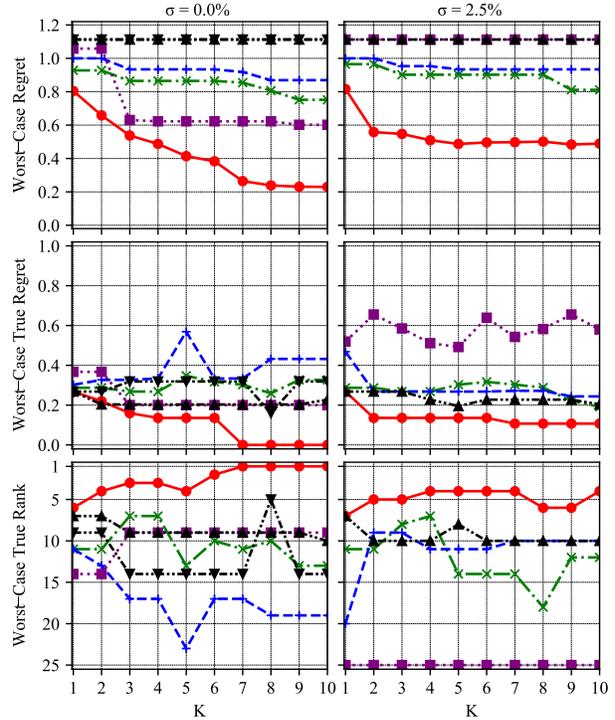


Table 6 Performance of our approach when the standard deviation σ of the inconsistencies used to select Γ in the optimization is different than the true standard deviation σ^* of the inconsistencies. Numbers are averages across values of K , σ , and σ^* . The detailed results can be found in Tables EC.3 (offline) and EC.4 (online).

Problem Type	Avg. Difference in Normalized Worst-Case Regret	Avg. Difference in Normalized Worst-Case True Regret	Avg. Difference in Worst-Case True Rank
Offline	0.03	0.03	2.44
Online	0.00	0.01	0.48

7. Conclusion

In this paper, motivated by the problem of making recommendations in high stakes settings where hard trade-offs need to be made, we proposed novel formulations for the offline and online max-min utility and min-max regret active preference elicitation problems. These take the form of robust optimization problems with decision-dependent information discovery. We studied the complexity of these problems and 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 Turk¹ 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<https://www.mturk.com>

Acknowledgments

P. Vayanos and E. Rice thank Dr. Iain DeJong for sharing the HMIS data used our numerical study. They also gratefully acknowledge support from Schmidt Futures and from the Hilton C. Foundation, the Homeless Policy Research Institute, and the Home for Good foundation under the “C.E.S. Triage Tool Research & Refinement” grant. P. Vayanos is also grateful for the support from the James H. Zumberge Faculty Research and Innovation Fund at the University of Southern California, from the U.S. Department of Transportation METTRANS University Transportation Center and National Center for Sustainable Transportation under Award No. DTRT13-G-UTC57, and from the National Science Foundation under CAREER award number 2046230. The authors are grateful to Marina Genchev, Sr. Manager for the Adult Coordinated Entry System at LAHSA, Erin Cox, former Policy Manager at LAHSA, and Maggie Potthoff, former 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 and Sina Aghaei and Nathan Jo for guidance on using and adapting their optimal classification trees and optimal prescriptive trees codes, respectively.

References

- Aghaei S, Gómez A, Vayanos P (2020) Strong optimal classification trees. Major revision at Operations Research, URL <https://arxiv.org/abs/2002.09142>.
- Agrawal S, Avadhanula V, Goyal V, Zeevi A (2019) MNL-bandit: A dynamic learning approach to assortment selection. *Operations Research* 67(5), ISSN 15265463, URL <http://dx.doi.org/10.1287/opre.2018.1832>.
- Araman VF, Caldentey R (2009) Dynamic Pricing for Nonperishable Products with Demand Learning. *Operations Research* 57:1169–1188.
- Armbruster B, Delage E (2015) Decision Making Under Uncertainty When Preference Information Is Incomplete. *Management Science* 61(1):111–128, URL <http://dx.doi.org/10.1287/mnsc.2014.2059>.
- Arora N, Huber J (2001) Improving parameter estimates and model prediction by aggregate customization in choice experiments. *Journal of Consumer Research* 28(2), ISSN 00935301, URL <http://dx.doi.org/10.1086/322902>.
- Arrow KJ (1963) *Social choice and individual values* (Yale University Press), 2nd edition.
- Auer P, Cesa-Bianchi N, Freund Y, Schapire RE (1995) Gambling in a rigged casino: the adversarial multi-armed bandit problem. *Annual Symposium on Foundations of Computer Science - Proceedings*, ISSN 02725428, URL <http://dx.doi.org/10.1109/sfcs.1995.492488>.

- Auer P, Cesa-Bianchi N, Freund Y, Schapire RE (2003) The nonstochastic multiarmed bandit problem. *SIAM Journal on Computing* 32(1), ISSN 00975397, URL <http://dx.doi.org/10.1137/S0097539701398375>.
- 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*.
- Badanidiyuru A, Langford J, Slivkins A (2014) Resourceful contextual bandits. *Journal of Machine Learning Research*, volume 35, ISSN 15337928.
- Ben-Tal A, El Ghaoui L, Nemirovski A (2009) *Robust Optimization*. Princeton Series in Applied Mathematics (Princeton University Press).
- Benade G, Nath S, Procaccia AD, Shah N (2017) Preference elicitation for participatory budgeting. *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (AAAI2017)*, 376–382.
- Bertsimas D, Brown D, Caramanis C (2010) Theory and applications of robust optimization. *SIAM Review* 53(3):464–501.
- Bertsimas D, Dunn J, Mundru N (2019) Optimal Prescriptive Trees. *INFORMS Journal on Optimization* 1(2), ISSN 2575-1484, URL <http://dx.doi.org/10.1287/ijoo.2018.0005>.
- 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, 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, Shtern S (2018) A Scalable Algorithm for Two-Stage Adaptive Linear Optimization. Technical report, Operations Research Center, MIT.
- Bertsimas D, Vayanos P (2017) Data-driven learning in dynamic pricing using adaptive robust optimization. Working paper, available on optimization online at http://www.optimization-online.org/DB_HTML/2014/10/4595.html.
- Besbes O, Zeevi A (2009) Dynamic pricing without knowing the demand function: risk bounds and near-optimal algorithms. *Operations Research* 57(6):1407–1420.
- Besbes O, Zeevi A (2012) On the (surprising) sufficiency of linear models for dynamic pricing with demand learning.
- Blasiak A, Khong J, Kee T (2020) CURATE.AI: Optimizing Personalized Medicine with Artificial Intelligence. URL <http://dx.doi.org/10.1177/2472630319890316>.
- 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.
- Broder J, Rusmevichientong P (2012) Dynamic pricing under a general parametric choice model. *Operations Research* 60(4):965–980.
- Brown DB, De Giorgi E, Sim M (2012) Aspirational preferences and their representation by risk measures. *Management Science* 58(11), ISSN 00251909, URL <http://dx.doi.org/10.1287/mnsc.1120.1537>.
- Brown DB, Sim M (2009) Satisficing measures for analysis of risky positions. *Management Science* 55(1), ISSN 00251909, URL <http://dx.doi.org/10.1287/mnsc.1080.0929>.
- Caro F, Gallien J (2007) Dynamic assortment with demand learning for seasonal consumer goods. *Management Science* 53(2), ISSN 00251909, URL <http://dx.doi.org/10.1287/mnsc.1060.0613>.
- 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.
- Chen W, Sim M (2009) Goal-driven optimization. *Operations Research* 57(2), ISSN 0030364X, URL <http://dx.doi.org/10.1287/opre.1080.0570>.
- Chen X, Krishnamurthy A, Wang Y (2019) Robust Dynamic Assortment Optimization in the Presence of Outlier Customers. *CoRR* abs/1910.0, URL <http://arxiv.org/abs/1910.04183>.
- Chen X, Wang Y, Zhou Y (2020) Dynamic assortment optimization with changing contextual information. *Journal of Machine Learning Research* 21, ISSN 15337928.
- Chen Y, Farias VF (2013) Simple policies for Dynamic pricing with imperfect Forecasts. *Operations Research* 61(3):612–624.
- Cheung WC, Simchi-Levi D (2017) Thompson Sampling for Online Personalized Assortment Optimization Problems with Multinomial Logit Choice Models. *SSRN Electronic Journal* ISSN 1556-5068, URL <http://dx.doi.org/10.2139/ssrn.3075658>.
- Cohen MC, Lobel I, Leme RP (2020) Feature-based dynamic pricing. *Management Science* 66(11), ISSN 15265501, URL <http://dx.doi.org/10.1287/mnsc.2019.3485>.
- 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.

- Dai J, Yang B, Guo C, Ding Z (2015) Personalized route recommendation using big trajectory data. *Proceedings - International Conference on Data Engineering*, volume 2015-May, ISSN 10844627, URL <http://dx.doi.org/10.1109/ICDE.2015.7113313>.
- Delage E, Li JYM (2018) Minimizing risk exposure when the choice of a risk measure is ambiguous. *Management Science* 64(1), ISSN 15265501, URL <http://dx.doi.org/10.1287/mnsc.2016.2593>.
- Delage E, Shaoyan G, Xu H (2018) Shortfall Risk Models When Information of Loss Function Is Incomplete. Technical report.
- den Boer AV, Zwart B (2013) Simultaneously learning and optimizing using controlled variance pricing.
- Dentcheva D, Ruszczyński A (2004) Optimization with stochastic dominance constraints. *SIAM Journal on Optimization* 14(2), ISSN 10526234, URL <http://dx.doi.org/10.1137/S1052623402420528>.
- Dudík M, Langford J, Li H (2011) Doubly robust policy evaluation and learning. *Proceedings of the 28th International Conference on Machine Learning, ICML 2011*.
- Dyer JS, Fishburn PC, Steuer RE, Wallenius J, Zionts S (1992) Multiple Criteria Decision Making, Multi-attribute Utility Theory: The Next Ten Years. *Management Science* 38(5):645–654.
- Feillet D, Gendreau M, Medaglia AL, Walteros JL (2010) A note on branch-and-cut-and-price. *Operations Research Letters* 38(5), ISSN 01676377, URL <http://dx.doi.org/10.1016/j.orl.2010.06.002>.
- Fischetti M, Vigo D (1997) A Branch-and-Cut Algorithm for the Resource-Constrained Minimum-Weight Arborescence Problem. *Networks* 29(1), ISSN 00283045, URL [http://dx.doi.org/10.1002/\(SICI\)1097-0037\(199701\)29:1<55::AID-NET6>3.0.CO;2-B](http://dx.doi.org/10.1002/(SICI)1097-0037(199701)29:1<55::AID-NET6>3.0.CO;2-B).
- Freedman R, Borg JS, Sinnott-Armstrong W, Dickerson JP, Conitzer V (2018) Adapting a kidney exchange algorithm to align with human values. *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society (AIES '18)*.
- 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.
- Gomez-Uribe CA, Hunt N (2016) The {N}etflix recommender system: Algorithms, business value, and innovation. *ACM Transactions on Management Information Systems (TMIS)* 6(4):13.
- Green PE, Srinivasan V (1978) Conjoint Analysis in Consumer Research: Issues and Outlook. *Journal of Consumer Research* 5(2), ISSN 0093-5301, URL <http://dx.doi.org/10.1086/208721>.

- 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>.
- Greene WWH (2012) *Econometric analysis 7th Ed*, volume 97.
- Gupta P, Goel A, Lin J, Sharma A, Wang D, Zadeh R (2013) WTF: The Who to Follow service at Twitter. *WWW 2013 - Proceedings of the 22nd International Conference on World Wide Web*.
- Gupta V, Grossmann IE (2011) Solution strategies for multistage stochastic programming with endogenous uncertainties. *Computers & Chemical Engineering* 35(11):2235–2247.
- Hardt M, Price E, Srebro N (2016) Equality of Opportunity in Supervised Learning. *ArXiv e-prints* .
- Harrison MJ, Keskin BN, Zeevi A (2012) Bayesian Dynamic Pricing Policies: Learning and Earning Under a Binary Prior Distribution. *Management Science* 58(3):570–586.
- Harrison MJ, Keskin NB, Zeevi A (2013) Dynamic pricing with an unknown linear demand model: asymptotic optimal semi-myopic policies. Technical report, The University of Chicago Booth School of Business.
- Haskell W, Huang W, Xu H (2018) Preference Elicitation and Robust Optimization with Multi-Attribute Quasi-Concave Choice Functions. Technical report.
- Hernán MA, Robins JM (2019) *Causal Inference: What If*. Chapman & Hall/CRC (Taylor & Francis), ISBN 9781420076165, URL https://books.google.com/books?id=%5C_KnHIAAACAAJ.
- Hill C, Hsu H, Holguin M, Morton M, Winetrobe H, Rice E (2021) An examination of housing interventions among youth experiencing homelessness: an investigation into racial/ethnic and sexual minority status. *Journal of Public Health* ISSN 1741-3842, URL <http://dx.doi.org/10.1093/pubmed/fdab295>.
- Hillier FS (2012) *Introduction To Operations Research* (McGraw-Hill Education (India) Pvt Limited), ISBN 9780071333467, URL <https://books.google.com/books?id=fsCjAgAAQBAJ>.
- Hsu HT, Hill C, Holguin M, Petry L, McElfresh D, Vayanos P, Morton M, Rice E (2021) Correlates of Housing Sustainability Among Youth Placed Into Permanent Supportive Housing and Rapid Re-Housing: A Survival Analysis. *Journal of Adolescent Health* 69(4), ISSN 18791972, URL <http://dx.doi.org/10.1016/j.jadohealth.2021.03.022>.
- Hu J, Bansal M, Mehrotra S (2018) Robust decision making using a general utility set. *European Journal of Operational Research* 269(2), ISSN 03772217, URL <http://dx.doi.org/10.1016/j.ejor.2018.02.018>.
- Hu J, Mehrotra S (2015) Robust decision making over a set of random targets or risk-averse utilities with an application to portfolio optimization. *IIE Transactions (Institute of Industrial Engineers)* 47(4), ISSN 15458830, URL <http://dx.doi.org/10.1080/0740817X.2014.919045>.
- Huber J, Zwerina K (1996) The Importance of Utility Balance in Efficient Choice Designs. *Journal of Marketing Research* 33(3), ISSN 00222437, URL <http://dx.doi.org/10.1177/002224379603300305>.

- Iancu DA, Trichakis N (2018) Monitoring With Limited Information. *Management Science* .
- Jo N, Aghaei S, Gomez A, Vayanos P (2021) Learning optimal prescriptive trees from observational data. Under review at Management Science, URL <https://arxiv.org/pdf/2108.13628.pdf>.
- Johnson FR, Lancsar E, Marshall D, Kilambi V, Mühlbacher A, Regier DA, Bresnahan BW, Kanninen B, Bridges JF (2013) Constructing experimental designs for discrete-choice experiments: Report of the ISPOR conjoint analysis experimental design good research practices task force. *Value in Health* 16(1), ISSN 10983015, URL <http://dx.doi.org/10.1016/j.jval.2012.08.2223>.
- Jones PW, Gittins JC (1989) Multi-Armed Bandit Allocation Indices. *The Journal of the Operational Research Society* 40(12), ISSN 01605682, URL <http://dx.doi.org/10.2307/2582928>.
- Jonsbråten TW (1998) *Optimization models for petroleum field exploitation*. Ph.D. thesis, Norwegian School 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.
- Kallus N (2017) Recursive partitioning for personalization using observational data. *34th International Conference on Machine Learning, ICML 2017*, volume 4.
- Kleinberg R, Leighton T (2003) The value of knowing a demand curve: bounds on regret for online posted-price auctions. *Proceedings on the 44th IEEE Symposium on Foundations of Computer Science*, 594–605.
- Kroes E, Sheldon R (1988) Stated preference methods: an introduction. *Journal of Transport Economics and Policy* 22(1), ISSN 0022-5258.
- Kube A, Das S, Fowler PJ (2019) Allocating Interventions Based on Predicted Outcomes: A Case Study on Homelessness Services. *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Kuhfeld WF (2010) Marketing research methods in SAS. . . . *Graphical Techniques*. Cary, NC, SAS-Institute TS-722 ISSN 00029270.
- Kuhfeld WF, Tobias RD, Garratt M (1994) Efficient Experimental Design with Marketing Research Applications. *Journal of Marketing Research* 31(4), ISSN 0022-2437, URL <http://dx.doi.org/10.1177/002224379403100408>.
- Ladeira JDM (2019) The algorithm and the flow: Netflix, machine learning and recommendation algorithms. *Intexto* (47), ISSN 1807-8583, URL <http://dx.doi.org/10.19132/1807-8583201947.166-184>.
- LAHSA (2017) Report and Recommendations of the Ad Hoc Committee on Women and Homelessness.
- LAHSA (2018) Report and Recommendations of the Ad Hoc Committee on Black People Experiencing Homelessness.
- Lam SW, Ng TS, Sim M, Song JH (2013) Multiple objectives satisficing under uncertainty. *Operations Research* 61(1), ISSN 0030364X, URL <http://dx.doi.org/10.1287/opre.1120.1132>.

- Li L, Chu W, Langford J, Schapire RE (2010) A contextual-bandit approach to personalized news article recommendation. *Proceedings of the 19th International Conference on World Wide Web, WWW '10*, URL <http://dx.doi.org/10.1145/1772690.1772758>.
- Löbel A (1998) Vehicle scheduling in public transit and Lagrangean pricing. *Management Science* 44(12 PART 1), ISSN 00251909, URL <http://dx.doi.org/10.1287/mnsc.44.12.1637>.
- Lu T, Pal D, Pal M (2010) Contextual Multi-Armed Bandits. Teh YW, Titterington M, eds., *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, volume 9 of *Proceedings of Machine Learning Research*, 485–492 (Chia Laguna Resort, Sardinia, Italy: PMLR), URL <https://proceedings.mlr.press/v9/lu10a.html>.
- Lykouris T, Mirrokni V, Leme RP (2018) Stochastic bandits robust to adversarial corruptions. *Proceedings of the Annual ACM Symposium on Theory of Computing*, ISSN 07378017, URL <http://dx.doi.org/10.1145/3188745.3188918>.
- Mahajan S, Van Ryzin G (2001) Stocking retail assortments under dynamic consumer substitution. *Operations Research* 49(3), ISSN 0030364X, URL <http://dx.doi.org/10.1287/opre.49.3.334.11210>.
- Mamer JW, McBride RD (2000) Decomposition-based pricing procedure for large-scale linear programs: an application to the linear multicommodity flow problem. *Management Science* 46(5), ISSN 00251909, URL <http://dx.doi.org/10.1287/mnsc.46.5.693.12042>.
- McFadden D (2001) Economic choices. *American Economic Review* 91(3), ISSN 00028282, URL <http://dx.doi.org/10.1257/aer.91.3.351>.
- McFadden DL (1974) Conditional Logit Analysis of Qualitative Choice Behavior. *Frontiers in Econometrics*.
- McSherry D (2002) Diversity-conscious retrieval. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 2416, ISSN 16113349, URL http://dx.doi.org/10.1007/3-540-46119-1_{_}17.
- Muter I, Birbil I, Bülbül K (2018) Benders decomposition and column-and-row generation for solving large-scale linear programs with column-dependent-rows. *European Journal of Operational Research* 264(1), ISSN 03772217, URL <http://dx.doi.org/10.1016/j.ejor.2017.06.044>.
- Muter I, Birbil SI, Bülbül K (2013) Simultaneous column-and-row generation for large-scale linear programs with column-dependent-rows. *Mathematical Programming* 142(1-2), ISSN 00255610, URL <http://dx.doi.org/10.1007/s10107-012-0561-8>.
- Natarajan K, Pachamanova D, Sim M (2007) Constructing Risk Measures from Uncertainty Sets. Working Paper.
- Nemhauser G, Wolsey L (1988) *Integer and Combinatorial Optimization*.
- Nemhauser GL, Wolsey LA, Fisher ML (1978) An analysis of approximations for maximizing submodular set functions—I. *Mathematical Programming* 14(1):265–294.

- Nohadani O, Roy A (2017) Robust optimization with time-dependent uncertainty in radiation therapy. *IIEE Transactions on Healthcare Systems Engineering* 7(2):81–92.
- Oh MH, Iyengar G (2019) Thompson sampling for multinomial logit contextual bandits. *Advances in Neural Information Processing Systems*, volume 32, ISSN 10495258.
- Orgcode (2014) Vulnerability Index (VI) and Service Prioritization Decision Assistance Tool (SPDAT). <http://org-beehivegroupcadev.nationbuilder.com/spdat>.
- Petry L, Hill C, Vayanos P, Rice E, Hsu HT, Morton M (2021) Associations Between the Vulnerability Index-Service Prioritization Decision Assistance Tool and Returns to Homelessness Among Single Adults in the United States. *Cityscape* 23(2):293–324.
- Rawls J (1991) Justice as Fairness: Political Not Metaphysical. *Equality and Liberty*, URL http://dx.doi.org/10.1007/978-1-349-21763-2{_}10.
- Reilly J, McCarthy K, McGinty L, Smyth B (2005) Incremental critiquing. *Knowledge-Based Systems*, volume 18, ISSN 09507051, URL <http://dx.doi.org/10.1016/j.knosys.2004.10.005>.
- Roth A, Slivkins A, Ullman J, Wu ZS (2020) Multidimensional Dynamic Pricing for Welfare Maximization. *ACM Transactions on Economics and Computation* 8(1), ISSN 21678383, URL <http://dx.doi.org/10.1145/3381527>.
- Roth A, Ullman J, Wu ZS (2016) Watch and learn: Optimizing from revealed preferences feedback. *Proceedings of the Annual ACM Symposium on Theory of Computing*, volume 19-21-June, ISSN 07378017, URL <http://dx.doi.org/10.1145/2897518.2897579>.
- Rudin C (2019) Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence* 1(5), ISSN 25225839, URL <http://dx.doi.org/10.1038/s42256-019-0048-x>.
- Rusmevichientong P, Shen ZJM, Shmoys DB (2010) Dynamic assortment optimization with a multinomial logit choice model and capacity constraint. *Operations Research* 58(6), ISSN 0030364X, URL <http://dx.doi.org/10.1287/opre.1100.0866>.
- Sadykov R, Vanderbeck F (2011) Column Generation for Extended Formulations. *Electronic Notes in Discrete Mathematics* 37:357–362, ISSN 1571-0653, URL <http://dx.doi.org/https://doi.org/10.1016/j.endm.2011.05.061>.
- Samuelson PA (1948) Consumption Theory in Terms of Revealed Preference. *Economica* 15(60), ISSN 00130427, URL <http://dx.doi.org/10.2307/2549561>.
- Sandholm T, Boutilier C (2006) Preference Elicitation in Combinatorial Auctions. Cramton P, Shoham Y, Steinberg R, eds., *Combinatorial Auctions*, 233–263 (MIT Press).
- Sauré D, Vielma JP (2019) Ellipsoidal methods for adaptive choice-based conjoint analysis. *Operations Research* 67(2).

- Sauré D, Zeevi A (2013) Optimal dynamic assortment planning with demand learning. *Manufacturing and Service Operations Management* 15(3), ISSN 15234614, URL <http://dx.doi.org/10.1287/msom.2013.0429>.
- Savage LJ (1951) The theory of statistical decision. *Journal of the American Statistical Association* 46(253):55–67.
- Stern H (1990) A continuum of paired comparisons models. *Biometrika* 77(2), ISSN 00063444, URL <http://dx.doi.org/10.1093/biomet/77.2.265>.
- Subrahmanyam S, Shoemaker R (1996) Developing optimal pricing and inventory policies for retailers who face Uncertain demand. *Journal of Retailing* 72(1):7–30.
- Syrkkanis V, Krishnamurthy A, Schapire RE (2016a) Efficient algorithms for adversarial contextual learning. *33rd International Conference on Machine Learning, ICML 2016*, volume 5.
- Syrkkanis V, Luo H, Krishnamurthy A, Schapire RE (2016b) Improved regret bounds for oracle-based adversarial contextual bandits. *Advances in Neural Information Processing Systems*, ISSN 10495258.
- 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.
- Toubia O, Hauser JR, Simester DI (2004) Polyhedral Methods for Adaptive Choice-Based Conjoint Analysis. *Journal of Marketing Research* 41(1):116–131.
- 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>.
- USRDS (2019) 2019 USRDS Annual Data Report: Epidemiology of kidney disease in the United States.
- Valério De Carvalho JM (1999) Exact solution of bin-packing problems using column generation and branch-and-bound. *Annals of Operations Research* 86, ISSN 02545330, URL <http://dx.doi.org/10.1023/A:1018952112615>.
- Van Ryzin G, Mahajan S (1999) On the relationship between inventory costs and variety benefits in retail assortments. *Management Science* 45(11), ISSN 00251909, URL <http://dx.doi.org/10.1287/mnsc.45.11.1496>.
- Vandenbergh L, Boyd S (2004) *Convex Optimization* (Cambridge University Press).
- Vayanos P, Georghiou A, Yu H (2019) Robust optimization with decision-dependent information discovery. R&R at Management Science, available online at <https://arxiv.org/abs/2004.08490>.
- 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.
- Viappiani P, Kroer C (2013) Robust optimization of recommendation sets with the maximin utility criterion. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 8176 LNAI, ISSN 16113349, URL http://dx.doi.org/10.1007/978-3-642-41575-3_{_}32.

-
- Von Neumann J (1928) Zur Theorie der Gesellschaftsspiele. *Mathematische Annalen* .
- 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.
- Wu H, Srikant R, Liu X, Jiang C (2015) Algorithms with logarithmic or sublinear regret for constrained contextual bandits. *Advances in Neural Information Processing Systems*, volume 2015-Janua, ISSN 10495258.
- Zeng B, Zhao L (2013) Solving two-stage robust optimization problems using a column-and-constraint generation method. *Operations Research Letters* 41(5):457–461.
- Zhao Z, Li H, Wang J, Kephart JO, Mattei N, Su H, Xia L (2018) A cost-effective framework for preference elicitation and aggregation. *34th Conference on Uncertainty in Artificial Intelligence 2018, UAI 2018*, volume 1.

E-Companion

EC.1. Illustration of the Uncertainty Set Update Procedure

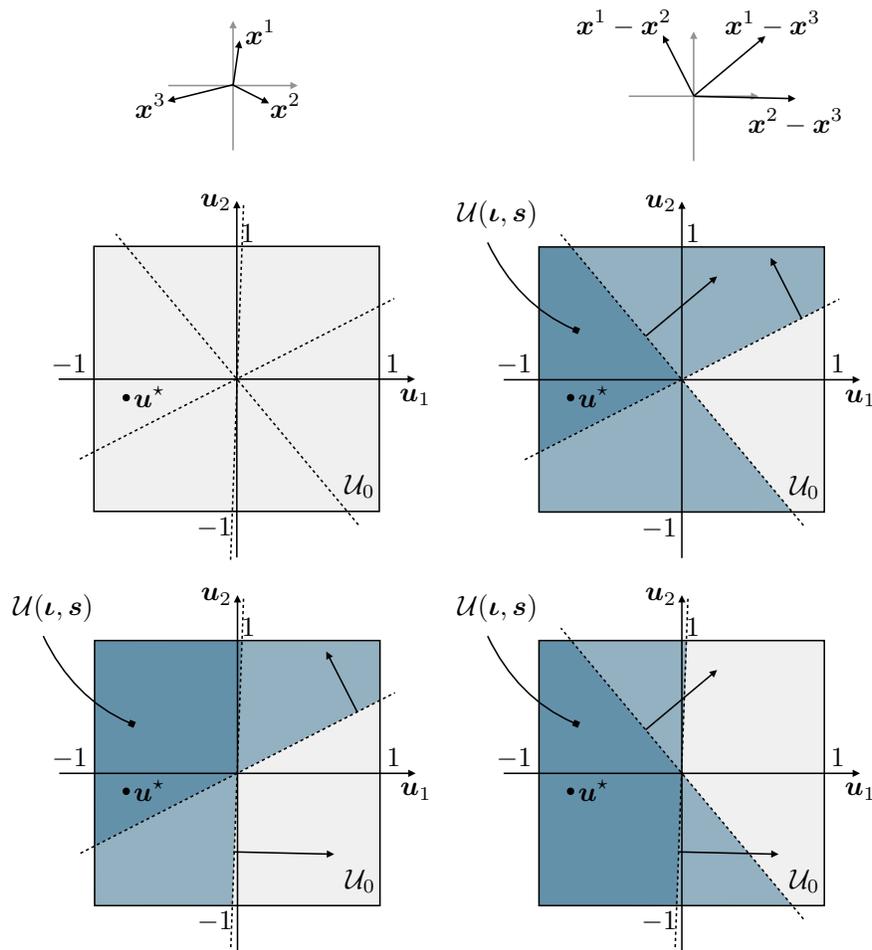


Figure EC.1 Illustration of the uncertainty set update procedure when there are $I = 3$ items, each with $J = 2$ features, $\Gamma = 0$, the “true” (but unknown) utility vector is \mathbf{u}^* , and the system is allowed to ask $K = 2$ questions before making a recommendation. The first row shows the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

EC.2. Generating Warm-Starts Satisfying Lexicographic Constraints

Algorithm 2 details the procedure for generating warm-starts proposed in Section 4.1.4 to ensure that the lexicographic constraints from Section 4.1.4 are satisfied.

Algorithm 2: An algorithm for building warm-starts

Inputs: A feasible solution $(\tau, \mathbf{v}, \mathbf{w}, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{x}, \bar{\mathbf{v}}, \bar{\mathbf{w}})$ to Problem (2) with K queries with objective τ ;

Output: A warm start $(\tilde{\mathbf{v}}, \tilde{\mathbf{w}}, \tilde{\mathbf{x}})$ to Problem (2) with $K + 1$ queries with objective $\geq \tau$;

Initialization:

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

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

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

end

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

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

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

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

end

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

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

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

end

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

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

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

end

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

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 i th element denotes the position of the i th output vector in the input;

EC.3. Greedy Solution Approach to Speed-Up Computations

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 all our experiments (unless explicitly stated otherwise). This heuristic is a variant of the warm-start approach, see Section 4.1.4 and is detailed in

Algorithm 3. This algorithm returns a feasible but potentially suboptimal solution to the MBLP counterpart of the offline preference elicitation problem to be solved. The idea of iteratively fixing some of the decision variables in the problem is not new and has a long standing history in operations research, see e.g., Nemhauser et al. (1978), Vandenberghe and Boyd (2004), Vayanos et al. (2019), Subrahmanyam and Shoemaker (1996).

Algorithm 3: Heuristic algorithm for solving Problem (\mathcal{QP}_u^K) (resp. (\mathcal{QP}_r^K)).

Inputs: Instance of Problem (\mathcal{QP}_u^K) (resp. (\mathcal{QP}_r^K));

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

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

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

 Solve the MBLP reformulation of Problem (\mathcal{QP}_u^1) , (resp. (\mathcal{QP}_r^1)) either directly or using the CCG algorithm;

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

else

 Solve the MBLP reformulation of Problem (\mathcal{QP}_u^κ) (resp. (\mathcal{QP}_r^κ)) with the added constraints that $\boldsymbol{\iota}^k = \boldsymbol{\iota}^{*,k}$ for all $k \in \{1, \dots, \kappa - 1\}$;

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

end

end

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

EC.4. Column-and-Constraint Generation for Min-Max Regret Active Preference Elicitation

We describe our column-and-constraint generation procedure using finite programs. Naturally, all problems solved would need to be converted to MBLPs first, using techniques similar to those employed in the proof of Theorem 5. We omit these conversions to streamline presentation.

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

minimize τ

subject to $\tau \in \mathbb{R}, \boldsymbol{\iota} \in \mathcal{C}^K, \boldsymbol{x}^s \in \mathcal{R}, \boldsymbol{s} \in \tilde{\mathcal{S}}^K$

$\boldsymbol{\alpha}^{(\boldsymbol{x}', \boldsymbol{s})} \in \mathbb{R}_-^K, \boldsymbol{\beta}^{(\boldsymbol{x}', \boldsymbol{s})} \in \mathbb{R}_-^M, \mu^{(\boldsymbol{x}', \boldsymbol{s})} \in \mathbb{R}_+$

$\tau \geq \boldsymbol{b}^\top \boldsymbol{\beta}^{(\boldsymbol{x}', \boldsymbol{s})} + \Gamma \mu^{(\boldsymbol{x}', \boldsymbol{s})}$

$\sum_{\kappa \in \mathcal{K}} \boldsymbol{s}_\kappa (\boldsymbol{x}^{\iota_\kappa^1} - \boldsymbol{x}^{\iota_\kappa^2}) \boldsymbol{\alpha}_\kappa^{(\boldsymbol{x}', \boldsymbol{s})} + \boldsymbol{B}^\top \boldsymbol{\beta}^{(\boldsymbol{x}', \boldsymbol{s})} = \boldsymbol{x}' - \boldsymbol{x}^s$

$\boldsymbol{\alpha}^{(\boldsymbol{x}', \boldsymbol{s})} + \mu^{(\boldsymbol{x}', \boldsymbol{s})} \boldsymbol{e} \geq \mathbf{0}$

$\left. \begin{array}{l} \forall \boldsymbol{s} \in \mathcal{S}', \\ \boldsymbol{x}' \in \mathcal{R}. \end{array} \right\} (\text{CCG}_r^{\text{main}}(\mathcal{S}'))$

This problem only involves a subset of the decision variables and constraints of Problem (7) (those indexed by $\mathbf{s} \in \mathcal{S}' \subseteq \tilde{\mathcal{S}}^K$). To identify scenarios that, given a solution $(\tau, \boldsymbol{\iota})$ to the relaxed main problem, must be added to $(\text{CCG}_r^{\text{main}}(\mathcal{S}'))$, we solve a *single* feasibility MBLP defined through

$$\begin{aligned}
& \text{maximize} && \theta \\
& \text{subject to} && \theta \in \mathbb{R}, \mathbf{s} \in \tilde{\mathcal{S}}^K \\
& && \mathbf{x}'^{\mathbf{x}} \in \mathcal{R}, \mathbf{u}^{\mathbf{x}} \in \mathcal{U}_0, \boldsymbol{\epsilon}^{\mathbf{x}} \in \mathcal{E}_\Gamma \quad \forall \mathbf{x} \in \mathcal{R} \\
& && \theta \leq (\mathbf{u}^{\mathbf{x}})^\top (\mathbf{x}'^{\mathbf{x}} - \mathbf{x}) \quad \forall \mathbf{x} \in \mathcal{R} \\
& && \left. \begin{aligned}
& (\mathbf{u}^{\mathbf{x}})^\top (\mathbf{x}^{\iota_1^{\kappa}} - \mathbf{x}^{\iota_2^{\kappa}}) + \boldsymbol{\epsilon}_\kappa^{\mathbf{x}} \geq M(\mathbf{s}_\kappa - 1) \\
& (\mathbf{u}^{\mathbf{x}})^\top (\mathbf{x}^{\iota_1^{\kappa}} - \mathbf{x}^{\iota_2^{\kappa}}) - \boldsymbol{\epsilon}_\kappa^{\mathbf{x}} \leq M(\mathbf{s}_\kappa + 1)
\end{aligned} \right\} \quad \forall \kappa \in \mathcal{K}, \mathbf{x} \in \mathcal{R}.
\end{aligned} \tag{CCG}_r^{\text{feas}}(\boldsymbol{\iota})$$

Problems $(\text{CCG}_r^{\text{main}}(\mathcal{S}'))$ and $(\text{CCG}_r^{\text{feas}}(\boldsymbol{\iota}))$ can be used in Algorithm 4 to solve Problem (\mathcal{QP}_r^K) . Convergence of the algorithm is guaranteed by the following theorem.

Theorem EC.1. *Algorithm 4 terminates in a final number of steps with a feasible solution to the regret averse active preference elicitation Problem (\mathcal{QP}_r^K) . The objective value attained by this solution is within δ of the optimal objective value of the problem.*

EC.5. Detail of Numerical Experiments on Synthetic Instances

We performed four set of experiments. In the first (resp. second) set of experiments, we evaluate the performance of our offline (resp. online) elicitation procedure for max-min utility (MMU) and min-max regret (MMR). In the third set of experiments, we investigate the computational performance of the CCG algorithm and the symmetry breaking constraints. In the last set of experiments, we investigate the relative performance of max-min utility and min-max regret solutions. Unless explicitly stated otherwise, we solve the offline elicitation problems (\mathcal{QP}_u^K) and (\mathcal{QP}_r^K) using the column-and-constraint generation procedures discussed in Sections 4.1.3 and 5.1.2, respectively, with tolerance $\delta = 1 \times 10^{-3}$. we augment all our formulations with the symmetry breaking constraints proposed in Section 4.1.4. To speed-up computation further, we also employ the greedy heuristic from Section EC.2. Without this conservative heuristic, the quality of the solutions obtained by our approach would be higher than those reported in these experiments.

For the first two sets of experiments, we investigate 7 settings of (I, J, σ) for each of MMU and MMR, where σ denotes the standard deviation of the inconsistencies. For each setting, we generate the features \mathbf{x}^i , $i \in \mathcal{I}$, uniformly at random from the J -dimensional sphere of radius ten. Throughout these experiments, we assume $\mathcal{U}^0 = [-1, 1]^J$ and set $\Gamma = 2\sigma\sqrt{K}\text{erf}^{-1}(2p - 1)$, where $p = 90\%$ is a confidence level and this definition is compatible with the elements of $\boldsymbol{\epsilon}$ being independent and normally distributed with standard deviation σ . Throughout our experiments, we use slightly

Algorithm 4: Column-and-Constraint Generation procedure for solving Problem $(\tilde{\mathcal{P}}_r^K)$.

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 ι^* from \mathcal{C}^K , near optimal in Problem $(\tilde{\mathcal{P}}_r^K)$ with associated objective θ ;

Initialization: $\iota^* \leftarrow \emptyset$; 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 $(\mathcal{CCG}_r^{\text{main}}(\mathcal{S}'))$. If it is solvable, let

$(\tau, \iota, \{\alpha^{(x',s)}, \beta^{(x',s)}\}_{x' \in \mathcal{R}, s \in \mathcal{S}'}, \{x^s\}_{s \in \mathcal{S}'})$ be an optimal solution. If it is unbounded, set $\tau = -\infty$ and let $\iota \in \mathcal{C}^K$ be such that $(\mathcal{CCG}_r^{\text{main}}(\mathcal{S}'))$ is unbounded when ι is fixed to that value;

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

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

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

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

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

end

end

Set $\iota^* \leftarrow \iota$;

Result: Collection of queries ι^* near-optimal in $(\tilde{\mathcal{P}}_r^K)$ with objective value θ .

larger datasets for MMU than for MMR. 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 MMR case.

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 between zero and one. Mathematically, the normalized worst-case utility associated with a choice of queries ι is calculated as

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

where u_{wc}^0 (resp. $u_{\text{wc}}^{\text{full}}$) denote the worst-case utility of the item recommended if no queries are asked (resp. if the utility is known). Thus, if after asking queries, recommendations are made optimally (to maximize worst case utility), 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 the analytic center of the uncertainty set as in Toubia et al. (2004)), then they may have a normalized worst-case utility that is negative. We can then interpret the normalized worst-case utility as a percentage relative to the full-information case (with 0% representing no information and 100% representing full-information). In the case of min-max regret, the normalized worst-case regret is calculated as

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

where r_{wc}^0 (resp. $r_{\text{wc}}^{\text{full}}$) denote the worst-case regret of the item recommended if no queries are asked (resp. if the utility is fully known). Thus, if after asking queries, recommendations are made to minimize worst-case regret, 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. We can then interpret the normalized worst-case regret as a percentage relative to the no-information case (with 100% representing no information and 0% representing full-information).

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 9.1.10 to compute the analytic centers of polyhedra as required by POLY, PROB, and ROB.

Offline Elicitation. In our offline experiments, we compare MMU (resp. MMR) to random elicitation (RAND) for $K \in \{1, \dots, 10\}$. For any fixed K , 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 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 $\boldsymbol{\iota}^*$, which we denote by $u_{\text{wc}}(\boldsymbol{\iota}^*)$ (resp. $r_{\text{wc}}(\boldsymbol{\iota}^*)$), by fixing $\boldsymbol{\iota} = \boldsymbol{\iota}^*$ in Problem (2) (resp. (7)). 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. The optimality and scalability results for MMU (resp. MMR) are provided in Figure EC.2 (resp. EC.3).

Online Elicitation. In the online setting, we compare our conservative solution approaches (MMU and MMR) to random elicitation, to the polyhedral method of Toubia et al. (2004) (POLY), to the probabilistic polyhedral method of Toubia et al. (2007) (PROB), to the robust approach of Bertsimas and O’Hair (2013) (ROB), and to the ellipsoidal method of Sauré and Vielma (2019) (ELL). For the online setting, we evaluate the performance of all approaches through simulation using 50 pairs $(\mathbf{u}, \boldsymbol{\epsilon})$. The utility vectors are drawn uniformly at random from the J -dimensional sphere of radius one. The inconsistencies are drawn from the normal distribution with standard deviation σ . For each method, we generate queries according to its elicitation approach, generate answers to the

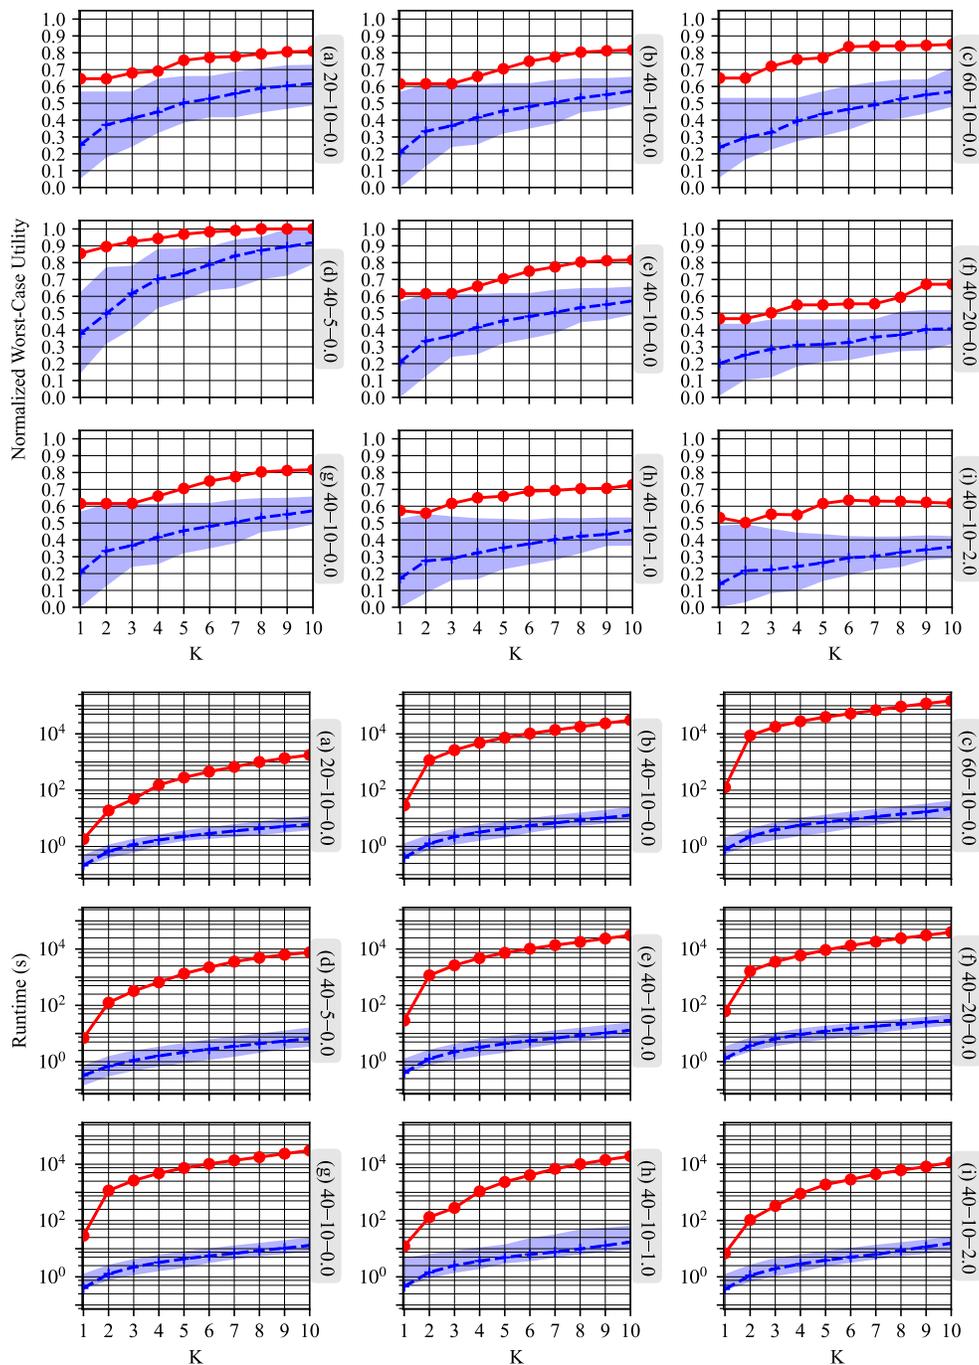


Figure EC.2 Optimality (top) and scalability (bottom) results for the offline max-min utility preference elicitation problem (QP_u^K) on synthetic data. Each label on the right of each facet corresponds to the characteristics of the instance solved ($I - J - \sigma$). Approach MMU is shown with red dots. The median performance of RAND across 50 sets of K random queries is shown with blue crosses. In particular the run time for RAND is the time it takes to evaluate the performance of one query. The blue shaded region shows the range of performance of RAND across the 50 sets of queries.

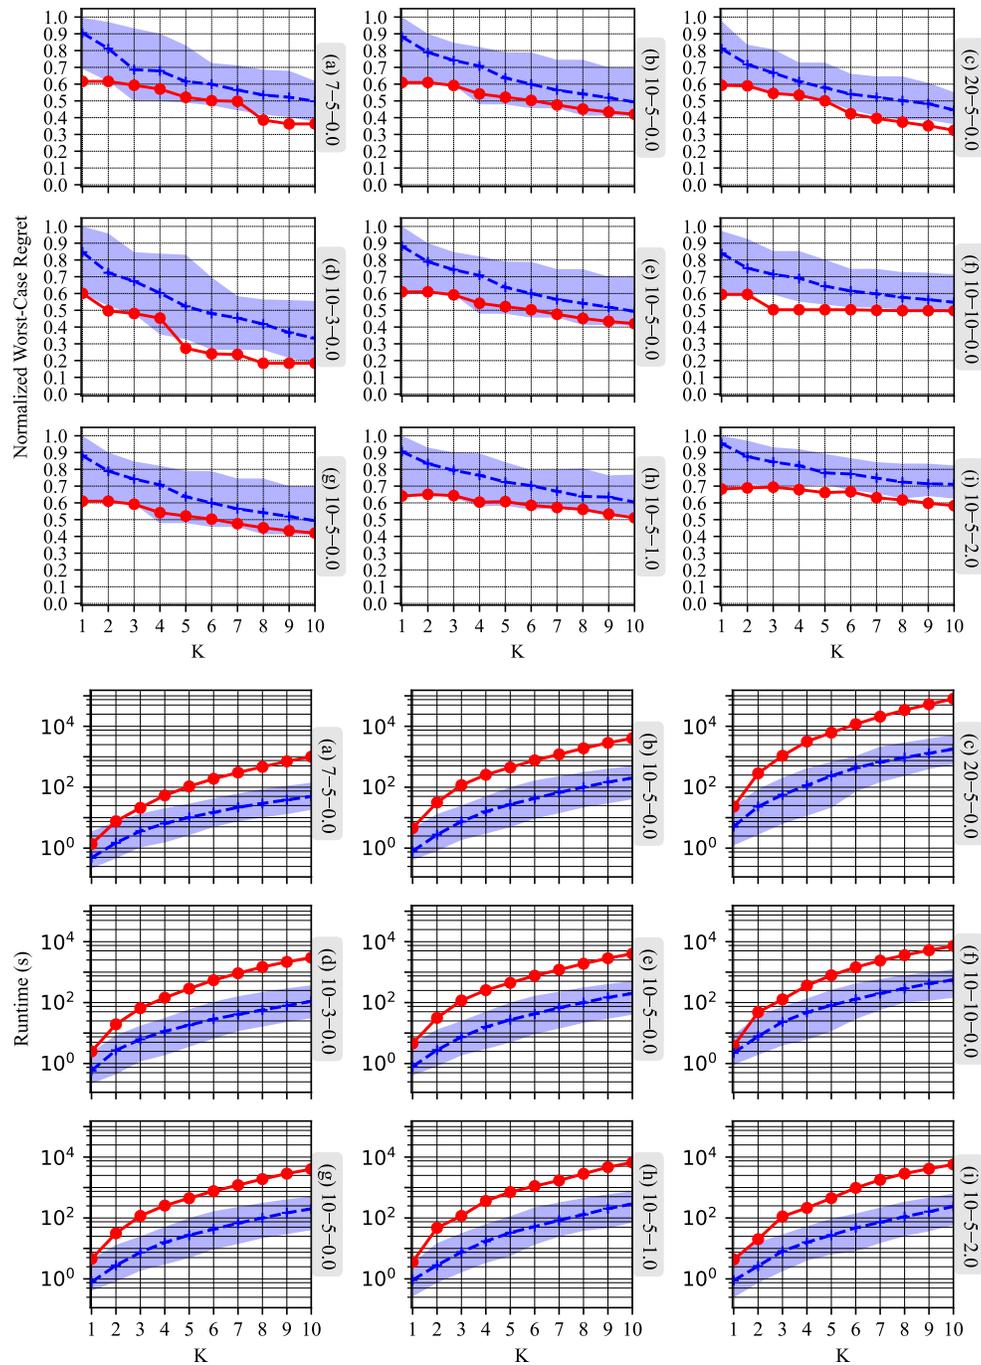


Figure EC.3 Optimality (top) and scalability (bottom) results for the offline min-max regret preference elicitation problem (QP_r^K) on synthetic data. The facet labels, graphs, shapes, lines, and colors have similar interpretation as in Figure EC.2.

queries using \mathbf{u} and ϵ , and subsequently make a recommendation according to the recommendation approach used by that method. We record the worst-case utility, worst-case regret, and rank of the recommended item. These are summarized in Figure EC.4 (resp. EC.5) for MMU (resp. MMR).

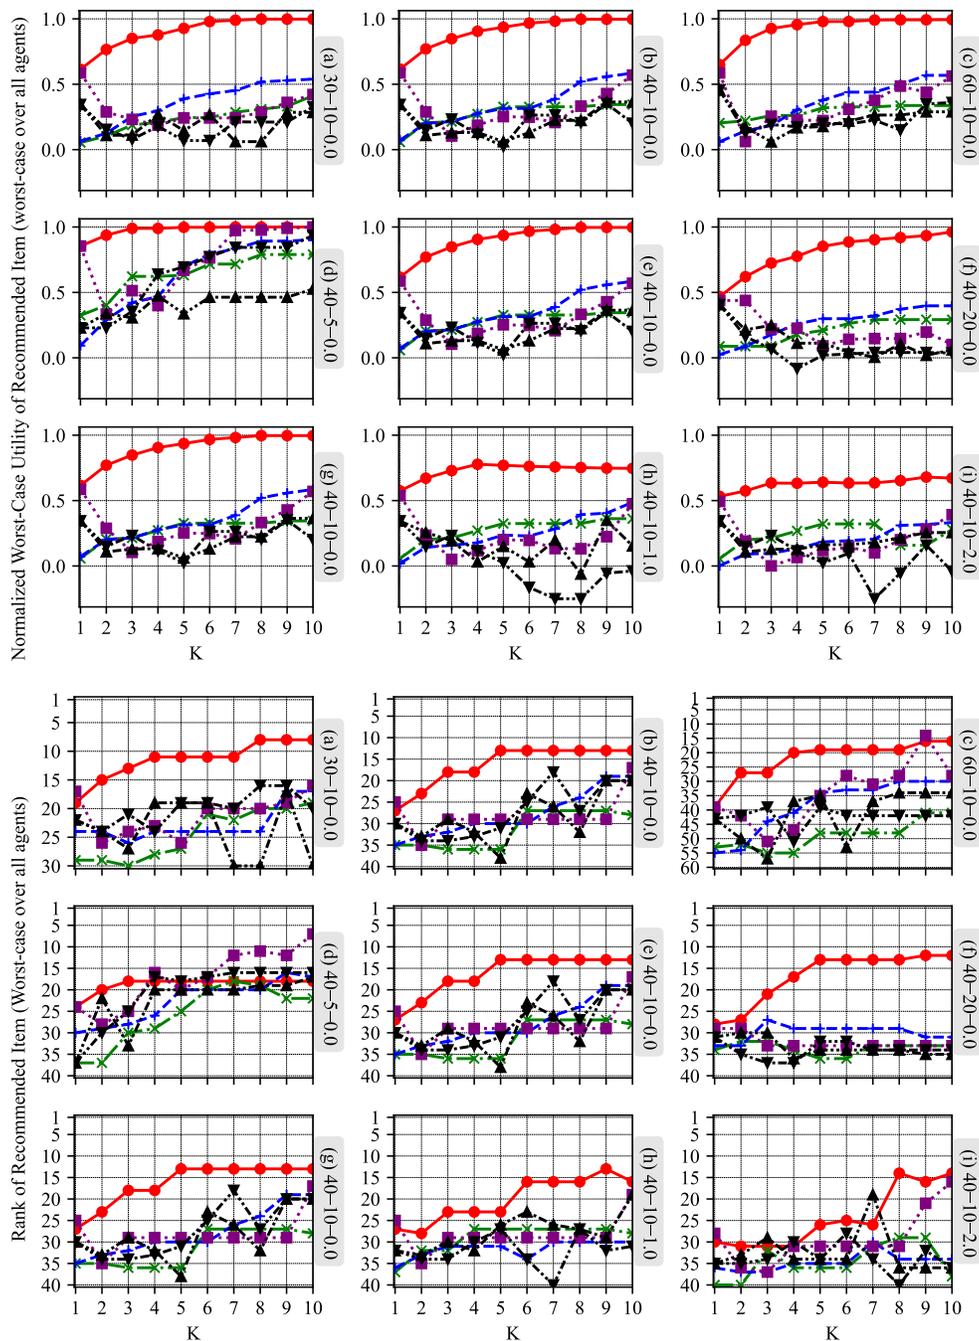


Figure EC.4 Worst-case utility (top) and worst-case rank (bottom) for the online max-min utility preference elicitation problem on synthetic data. The facet labels have the same interpretation as in Figure EC.2. The median performance of MMU (resp. RAND, POLY, PROB, ROB, ELL) across 50 random utility vectors \mathbf{u} and inconsistencies ϵ is shown with red dots (resp. blue plus, black triangles, black inverted triangles, green crosses, and purple squares).

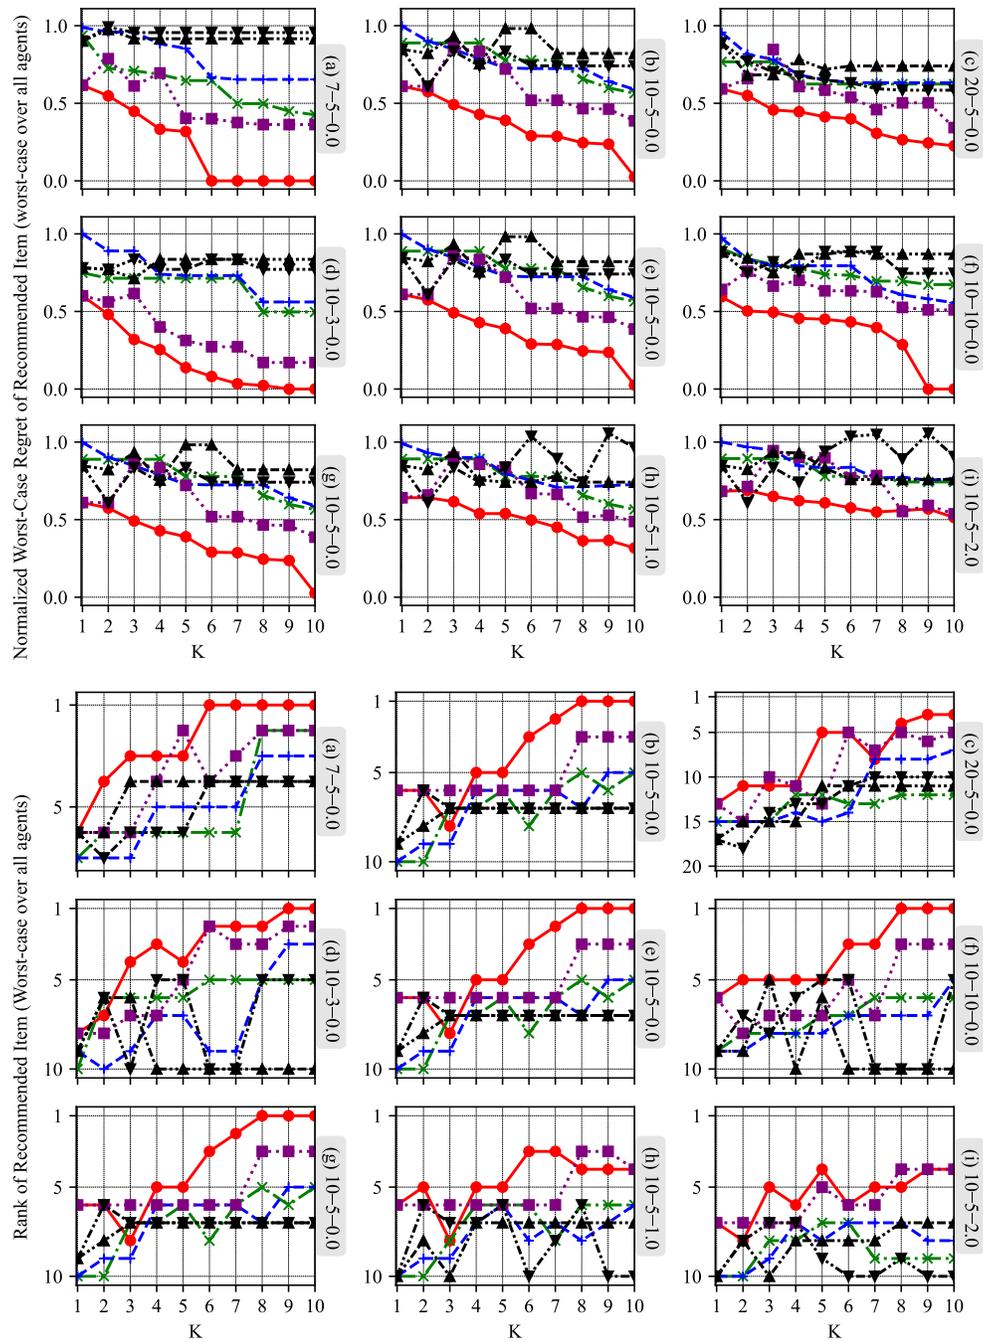


Figure EC.5 Worst-case utility (top) and worst-case rank (bottom) for the online min-max-regret preference elicitation problem on synthetic data. The facet labels and colors have the same interpretation as in Figure EC.4 with the exception that red dots correspond to MMR.

Comparison Between Max-Min Utility & Min-Max Regret. 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 third set of experiments, we investigate whether there are benefits in employing the min-max

Table EC.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 \mathbf{u}_u^* and min-max regret queries \mathbf{u}_r^* are computed by solving the MBLP reformulations of Problems (\mathcal{QP}_u^K) and (\mathcal{QP}_r^K) , 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 $(u_{\text{nwc}}(\mathbf{u}_u^*) - u_{\text{nwc}}(\mathbf{u}_r^*))$. 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 $(r_{\text{nwc}}(\mathbf{u}_u^*) - r_{\text{nwc}}(\mathbf{u}_r^*))$. All approaches were given a one hour time limit.

K	Max-Min Utility Solution		Min-Max Regret Solution		Decrease in Normalized Worst-Case Utility (p.p.)	Decrease in Normalized Worst-Case Regret (p.p.)
	Normalized	Normalized	Normalized	Normalized		
	Worst-Case Utility	Worst-Case Regret	Worst-Case Utility	Worst-Case Regret		
2	0.780	0.642	0.689	0.594	9.1	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	5.1
10	0.912	0.510	0.844	0.495	6.8	1.5

Table EC.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)						
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

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{QP}_u^K) and (\mathcal{QP}_r^K) , respectively, on a synthetic dataset with $I = 10$ items and $J = 10$ features ($\Gamma = 0$). The results are summarized in Table EC.1.

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{QP}_u^K) based on three synthetic datasets using four different approaches: the MBLP problem (2), the MBLP problem (2) augmented with symmetry breaking constraints (see Section 4.1.4), the column-and-constraint generation (CCG) approach from Algorithm 1 augmented with symmetry breaking constraints, and the CCG-based conservative heuristic approach, see Section EC.3. The results are provided in Table EC.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.

EC.6. Detail of Numerical Experiments on Real Data

In Tables EC.3 and EC.4, we provide the full detail of our results on evaluating robustness of our model to the choice of Γ in the real dataset.

Table EC.3 Performance of the offline MMR approach on the real data when the standard deviation σ of the inconsistencies used to select Γ in the optimization is different than the true standard deviation σ^* of the inconsistencies.

Metric	K	$\sigma = 0.0\%$		$\sigma = 2.5\%$			$\sigma = 5.0\%$		
		$\sigma^* = 0.0\%$	$\sigma^* = 0.5\%$	$\sigma^* = 2.0\%$	$\sigma^* = 2.5\%$	$\sigma^* = 3.0\%$	$\sigma^* = 4.5\%$	$\sigma^* = 5.0\%$	$\sigma^* = 5.5\%$
Worst-Case Regret	1	0.804	0.804	0.898	0.898	0.898	0.937	0.937	0.937
	2	0.716	0.716	0.810	0.810	0.810	0.934	0.934	0.934
	3	0.671	0.671	0.780	0.780	0.780	0.934	0.934	0.934
	4	0.667	0.642	0.757	0.709	0.757	0.929	0.929	0.934
	5	0.658	0.641	0.721	0.630	0.712	0.925	0.816	0.926
	6	0.640	0.632	0.699	0.608	0.686	0.814	0.804	0.804
	7	0.597	0.599	0.689	0.607	0.670	0.804	0.797	0.779
	8	0.547	0.551	0.670	0.581	0.651	0.787	0.769	0.746
	9	0.539	0.538	0.670	0.581	0.635	0.773	0.727	0.686
	10	0.538	0.533	0.652	0.578	0.612	0.758	0.689	0.667
Worst-Case True Regret	1	0.305	0.305	0.367	0.402	0.367	0.359	0.359	0.323
	2	0.305	0.305	0.305	0.305	0.305	0.360	0.321	0.311
	3	0.305	0.305	0.340	0.310	0.340	0.360	0.305	0.334
	4	0.305	0.305	0.391	0.294	0.340	0.360	0.360	0.323
	5	0.305	0.305	0.391	0.244	0.340	0.347	0.351	0.315
	6	0.305	0.351	0.341	0.340	0.334	0.347	0.351	0.318
	7	0.341	0.351	0.341	0.319	0.340	0.351	0.351	0.318
	8	0.341	0.217	0.329	0.301	0.334	0.358	0.351	0.315
	9	0.341	0.217	0.329	0.301	0.334	0.365	0.351	0.223
	10	0.341	0.224	0.329	0.340	0.334	0.351	0.351	0.264
Worst-Case True Rank	1	11	11	21	23	21	20	18	18
	2	11	11	11	11	11	15	15	15
	3	11	11	18	18	19	14	16	14
	4	11	11	20	16	13	16	20	20
	5	11	9	20	8	11	17	18	17
	6	11	17	19	12	15	15	17	20
	7	13	13	19	12	15	15	20	20
	8	13	8	15	12	22	20	17	17
	9	13	8	15	12	15	15	17	19
	10	12	8	14	11	12	15	15	16

Table EC.4 Performance of the online MMR approach on the real data when the standard deviation σ of the inconsistencies used to select Γ in the optimization is different than the true standard deviation σ^* of the inconsistencies.

Metric	K	$\sigma = 0.0\%$		$\sigma = 2.5\%$			$\sigma = 5.0\%$		
		$\sigma^* = 0.0\%$	$\sigma^* = 0.5\%$	$\sigma^* = 2.0\%$	$\sigma^* = 2.5\%$	$\sigma^* = 3.0\%$	$\sigma^* = 4.5\%$	$\sigma^* = 5.0\%$	$\sigma^* = 5.5\%$
Worst-Case Regret	1	0.804	0.804	0.816	0.816	0.816	0.816	0.816	0.853
	2	0.658	0.658	0.558	0.558	0.558	0.635	0.635	0.635
	3	0.537	0.537	0.547	0.547	0.547	0.607	0.607	0.607
	4	0.487	0.487	0.509	0.509	0.509	0.616	0.616	0.616
	5	0.414	0.414	0.487	0.487	0.487	0.635	0.635	0.635
	6	0.383	0.383	0.496	0.496	0.496	0.652	0.652	0.652
	7	0.264	0.264	0.497	0.497	0.497	0.661	0.661	0.661
	8	0.238	0.238	0.501	0.501	0.501	0.622	0.622	0.622
	9	0.231	0.231	0.483	0.483	0.483	0.558	0.558	0.558
	10	0.229	0.229	0.489	0.489	0.489	0.566	0.566	0.566
Worst-Case True Regret	1	0.268	0.268	0.268	0.268	0.268	0.268	0.268	0.245
	2	0.219	0.219	0.135	0.135	0.135	0.135	0.135	0.121
	3	0.159	0.159	0.135	0.135	0.135	0.135	0.184	0.165
	4	0.135	0.135	0.135	0.135	0.135	0.135	0.159	0.143
	5	0.135	0.135	0.135	0.135	0.135	0.135	0.135	0.121
	6	0.135	0.135	0.135	0.135	0.135	0.176	0.176	0.158
	7	0	0.008	0.107	0.107	0.107	0.135	0.135	0.121
	8	0	0.008	0.107	0.107	0.107	0.135	0.135	0.121
	9	0	0.008	0.107	0.107	0.107	0.135	0.135	0.121
	10	0	0.008	0.107	0.107	0.107	0.135	0.135	0.121
Worst-Case True Rank	1	6	6	7	7	7	7	7	15
	2	4	4	5	5	5	5	5	5
	3	3	3	5	5	5	5	6	6
	4	3	3	4	4	4	5	5	5
	5	4	6	4	4	4	5	5	5
	6	2	3	4	4	4	5	5	5
	7	1	3	4	4	4	3	3	3
	8	1	3	4	6	6	5	5	5
	9	1	3	4	6	6	7	7	7
	10	1	3	4	4	4	7	7	7

EC.7. Proofs of Statements in Section 4

The following proposition is needed in the proof of Lemma 1. It shows that we can eliminate the dependence of $\mathcal{S}_\Gamma(\boldsymbol{\iota})$ on $\boldsymbol{\iota}$ in Problem (\mathcal{QP}_u^K) .

Proposition EC.1. *Problem (\mathcal{QP}_u^K) is equivalent to*

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

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

$$\min_{\mathbf{s} \in \mathcal{S}^K \setminus \{\mathcal{S}_\Gamma(\boldsymbol{\iota})\}} \max_{\mathbf{x} \in \mathcal{R}} \min_{\mathbf{u} \in \mathcal{U}_\Gamma(\boldsymbol{\iota}, \mathbf{s})} \mathbf{u}^\top \mathbf{x} = +\infty.$$

Therefore,

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

This implies that

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

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

Proof of Lemma 1 We begin by showing that Problem (EC.1) is equivalent to

$$\text{maximize}_{\boldsymbol{\iota} \in \mathcal{C}^K} \min_{\mathbf{s} \in \mathcal{S}^K} \max_{\mathbf{x} \in \mathcal{R}} \min_{\mathbf{u} \in \bar{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \mathbf{s})} \mathbf{u}^\top \mathbf{x}, \quad (\text{EC.2})$$

where

$$\bar{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \mathbf{s}) := \left\{ \mathbf{u} \in \mathcal{U}^0 : \exists \boldsymbol{\epsilon} \in \mathcal{E}_\Gamma \text{ such that } \begin{cases} \mathbf{u}^\top (\mathbf{x}^{\boldsymbol{\iota}_1^\kappa} - \mathbf{x}^{\boldsymbol{\iota}_2^\kappa}) \geq -\boldsymbol{\epsilon}_\kappa & \forall \kappa \in \mathcal{K} : \mathbf{s}_\kappa = 1 \\ |\mathbf{u}^\top (\mathbf{x}^{\boldsymbol{\iota}_1^\kappa} - \mathbf{x}^{\boldsymbol{\iota}_2^\kappa})| \leq \boldsymbol{\epsilon}_\kappa & \forall \kappa \in \mathcal{K} : \mathbf{s}_\kappa = 0 \\ \mathbf{u}^\top (\mathbf{x}^{\boldsymbol{\iota}_1^\kappa} - \mathbf{x}^{\boldsymbol{\iota}_2^\kappa}) \leq \boldsymbol{\epsilon}_\kappa & \forall \kappa \in \mathcal{K} : \mathbf{s}_\kappa = -1 \end{cases} \right\}.$$

To this end, fix any $\boldsymbol{\iota} \in \mathcal{C}^K$. We show that

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

and

$$\min_{\mathbf{s} \in \mathcal{S}^K} \max_{\mathbf{x} \in \mathcal{R}} \min_{\mathbf{u} \in \bar{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \mathbf{s})} \mathbf{u}^\top \mathbf{x} \quad (\text{EC.4})$$

attain the same optimal value.

Since $\mathcal{U}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s}) \subseteq \overline{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s})$, the optimal objective value of Problem (EC.4) constitutes a lower bound on the optimal objective value of Problem (EC.3).

For the converse part, observe that, by virtue of the linearity of $\mathbf{u}^\top \mathbf{x}$ and the boundedness of $\mathcal{U}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s})$, Problems (EC.3) and (EC.4) have the same objective value for all $\boldsymbol{s} \in \mathcal{S}^K$ such that $\mathcal{U}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s}) \neq \emptyset$ since, in such cases, $\overline{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s}) = \text{cl}(\mathcal{U}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s}))$. Let \boldsymbol{s}^* be optimal in Problem (EC.4). If $\mathcal{U}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s}^*) \neq \emptyset$, the proof is complete. Suppose instead that $\mathcal{U}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s}^*) = \emptyset$. Since $\mathcal{U}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s}^*) = \emptyset$ while $\overline{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s}^*) \neq \emptyset$, this implies that there exist implied equalities in the set $\overline{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s}^*)$, which may only happen if $\Gamma = 0$. In particular, since \mathcal{U}^0 is full-dimensional, see Assumption 1, the implied equalities are all associated with response constraints. Let us collect the indices κ of all those implied equalities in the set $\mathcal{K}_=$, i.e., let

$$\mathcal{K}_= := \left\{ \kappa \in \mathcal{K} : \mathbf{s}_\kappa \in \{-1, 1\}, \mathbf{u}^\top (\mathbf{x}^{\iota_1^\kappa} - \mathbf{x}^{\iota_2^\kappa}) = 0 \quad \forall \mathbf{u} \in \overline{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s}^*) \right\}.$$

Define

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

This definition of \mathbf{s}' ensures that equality constraints that are implicit in the set $\overline{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s}^*)$ are made explicit in the set $\overline{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s}')$. By construction, $\overline{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s}') = \overline{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s}^*)$ and thus \boldsymbol{s}' is optimal in Problem (EC.4). Moreover, since $\overline{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s}')$ is non-empty, it follows by Proposition 2.3 in Nemhauser and Wolsey (1988) that $\mathcal{U}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s}')$ is non-empty. Thus, $\overline{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s}') = \text{cl}(\mathcal{U}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s}'))$, implying that the objective value attained by \boldsymbol{s}' in Problem (EC.3) equals the optimal objective value of Problem (EC.4). This in turn yields the required result that the optimal objective value of Problem (EC.4) constitutes an upper bound on the optimal objective value of Problem (EC.3).

Thus, the optimal values of Problems (EC.3) and (EC.4) are equal implying that Problems (EC.1) and (EC.2) are equivalent.

It now suffices to show that Problems (EC.2) and $(\widetilde{\mathcal{QP}}_u^K)$ are equivalent. Fix any $\boldsymbol{\iota} \in \mathcal{C}^K$. Let \boldsymbol{s}^* be optimal in

$$\min_{\boldsymbol{s} \in \mathcal{S}^K} \max_{\mathbf{x} \in \mathcal{R}} \min_{\mathbf{u} \in \overline{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s})} \mathbf{u}^\top \mathbf{x} \tag{EC.5}$$

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

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

Then, $\overline{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s}') \supseteq \overline{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s}^*)$, implying that the objective value attained by \boldsymbol{s}' in Problem (EC.5) is no greater than the objective value attained by \boldsymbol{s}^* . Thus, \boldsymbol{s}' is also optimal in (EC.5). Since $\boldsymbol{s}' \in \widetilde{\mathcal{S}}^K$, this implies that we can always restrict our search for an optimal solution to Problem (EC.5) to the set $\widetilde{\mathcal{S}}^K$. As the choice of $\boldsymbol{\iota} \in \mathcal{C}^K$ was arbitrary and noting that $\overline{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s}) = \widetilde{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s})$ for all $\boldsymbol{s} \in \widetilde{\mathcal{S}}^K$ and $\boldsymbol{\iota} \in \mathcal{C}^K$, this concludes the proof. \square

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

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

Since the choices of $\Gamma \in \mathbb{R}_+$, $K \in \mathbb{N}$, $\boldsymbol{\iota} \in \mathcal{C}^K$, and $\boldsymbol{s} \in \tilde{\mathcal{S}}^K$ such that $\tilde{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s}) \neq \emptyset$ were arbitrary, it holds that (under the premise of the theorem statement) Problem $(\widetilde{\mathcal{QP}}_u^K)$ is equivalent to

$$\text{maximize}_{\boldsymbol{\iota} \in \mathcal{C}^K} \min_{\boldsymbol{s} \in \tilde{\mathcal{S}}^K} \min_{\boldsymbol{u} \in \tilde{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s})} \max_{\boldsymbol{x} \in \mathcal{R}} \boldsymbol{u}^\top \boldsymbol{x}. \quad (\text{EC.6})$$

Fix $\Gamma \in \mathbb{R}_+$ and $K \in \mathbb{N}$. Then,

$$\begin{aligned} & \max_{\boldsymbol{\iota} \in \mathcal{C}^K} \min_{\boldsymbol{s} \in \tilde{\mathcal{S}}^K} \min_{\boldsymbol{u} \in \tilde{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s})} \max_{\boldsymbol{x} \in \mathcal{R}} \boldsymbol{u}^\top \boldsymbol{x} \\ &= \max_{\boldsymbol{\iota} \in \mathcal{C}^K} \min_{\substack{\boldsymbol{u} \in \bigcup_{\boldsymbol{s} \in \tilde{\mathcal{S}}^K} \\ \tilde{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s})}} \max_{\boldsymbol{x} \in \mathcal{R}} \boldsymbol{u}^\top \boldsymbol{x} \\ &= \max_{\boldsymbol{\iota} \in \mathcal{C}^K} \min_{\boldsymbol{u} \in \mathcal{U}^0} \max_{\boldsymbol{x} \in \mathcal{R}} \boldsymbol{u}^\top \boldsymbol{x} \\ &= \min_{\boldsymbol{u} \in \mathcal{U}^0} \max_{\boldsymbol{x} \in \mathcal{R}} \boldsymbol{u}^\top \boldsymbol{x} \\ &= \max_{\boldsymbol{x} \in \mathcal{R}} \min_{\boldsymbol{u} \in \mathcal{U}^0} \boldsymbol{u}^\top \boldsymbol{x}, \end{aligned} \quad (\text{EC.7})$$

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

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

and the second equality follows by definitions of $\tilde{\mathcal{S}}^K$ and $\tilde{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s})$. 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).

<p style="margin: 0;">PARTITION.</p> <p>Instance. Given a set \mathcal{A} of elements $\mathcal{A} := \{1, \dots, n\}$ with associated positive integer weights $\boldsymbol{w}_i \in \mathbb{N}_+$, $i \in \mathcal{A}$, such that $\sum_{i \in \mathcal{A}} \boldsymbol{w}_i = 2W$.</p> <p>Question. Does there exist a partition of \mathcal{A} into two subsets, \mathcal{X} and $\mathcal{A} \setminus \mathcal{X}$, such that $\sum_{i \in \mathcal{X}} \boldsymbol{w}_i = \sum_{i \in \mathcal{A} \setminus \mathcal{X}} \boldsymbol{w}_i = W$?</p>

We aim to reduce the partition problem to evaluating the objective function of an instance of Problem $(\widetilde{\mathcal{QP}}_u^K)$. To this end, fix an instance (n, \mathbf{w}, W) of the partition problem. Set $\Gamma = 0$, $J = n + 1$, $K = n$, and $\mathcal{U}^0 = \{\mathbf{u} \in [0, 1]^J : \mathbf{u}_J = 0.5\}$. 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 $\mathbf{s} \in \widetilde{\mathcal{S}}^K$, we have

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

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

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

Then, we have

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

Now, we claim that we are given a “yes” instance of PARTITION if and only if the objective value of $\boldsymbol{\iota}$ in the constructed instance of Problem $(\widetilde{\mathcal{QP}}_u^K)$ is 0. To this end, note that the objective value of $\boldsymbol{\iota}$ is given by

$$\text{minimize}_{\mathbf{s} \in \widetilde{\mathcal{S}}^K} Z(\mathbf{s}) \tag{EC.8}$$

and is lower bounded by 0. If there exists $\mathcal{X} \subset \mathcal{A}$ such that $\sum_{i \in \mathcal{X}} \mathbf{w}_i = W$, then the solution $\mathbf{s} \in \widetilde{\mathcal{S}}^K$ defined through $\mathbf{s}_\kappa = 1$ if $\kappa \in \mathcal{X}$, $\mathbf{s}_\kappa = -1$ else, $\kappa \in \mathcal{K}$, attains an objective value of zero in Problem (EC.8). Conversely, if the optimal objective value of Problem (EC.8) is zero, then the set $\mathcal{X} := \{\kappa \in \mathcal{A} : \mathbf{s}_\kappa = 1\}$ is such that $\sum_{i \in \mathcal{X}} \mathbf{w}_i = W$ and the claim follows. This concludes the proof of the second item.

Both items are thus proved. \square

The following observation is needed in the proof of Lemma 2.

Observation EC.1. Problem $(\widetilde{\mathcal{QP}}_u^K)$ is equivalent to the following max-min problem

$$\text{maximize}_{\boldsymbol{\iota} \in \mathcal{C}^K} \quad \max_{\substack{\mathbf{x}^s \in \mathcal{R}: \\ \mathbf{s} \in \widetilde{\mathcal{S}}^K}} \quad \min_{\mathbf{s} \in \widetilde{\mathcal{S}}^K} \quad \min_{\mathbf{u} \in \widetilde{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \mathbf{s})} \quad \mathbf{u}^\top \mathbf{x}^s, \quad (\text{EC.9})$$

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

Proof of Observation EC.1 Fix $\boldsymbol{\iota} \in \mathcal{C}^K$. Then,

$$\begin{aligned} & \min_{\mathbf{s} \in \widetilde{\mathcal{S}}^K} \quad \max_{\mathbf{x} \in \mathcal{R}} \quad \min_{\mathbf{u} \in \widetilde{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \mathbf{s})} \quad \mathbf{u}^\top \mathbf{x} \\ = & \min_{\mathbf{s} \in \widetilde{\mathcal{S}}^K} \left\{ \max_{\mathbf{x}^s \in \mathcal{R}} \quad \min_{\mathbf{u} \in \widetilde{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \mathbf{s})} \quad \mathbf{u}^\top \mathbf{x}^s \right\} \\ = & \max_{\substack{\mathbf{x}^s \in \mathcal{R}: \\ \mathbf{s} \in \widetilde{\mathcal{S}}^K}} \quad \min_{\mathbf{s} \in \widetilde{\mathcal{S}}^K} \left\{ \min_{\mathbf{u} \in \widetilde{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \mathbf{s})} \quad \mathbf{u}^\top \mathbf{x}^s \right\} \\ = & \max_{\substack{\mathbf{x}^s \in \mathcal{R}: \\ \mathbf{s} \in \widetilde{\mathcal{S}}^K}} \quad \min_{\mathbf{s} \in \widetilde{\mathcal{S}}^K} \quad \min_{\mathbf{u} \in \widetilde{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \mathbf{s})} \quad \mathbf{u}^\top \mathbf{x}^s, \end{aligned} \quad (\text{EC.10})$$

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 (EC.9) can be written in epigraph form equivalently as

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

The problem above is in turn equivalent to

$$\begin{aligned} & \text{maximize} \quad \tau \\ & \text{subject to} \quad \tau \in \mathbb{R}, \boldsymbol{\iota} \in \mathcal{C}^K, \mathbf{x}^s \in \mathcal{R}, \mathbf{s} \in \widetilde{\mathcal{S}}^K \\ & \quad \quad \quad \tau \leq \min_{\mathbf{u} \in \widetilde{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \mathbf{s})} \quad \mathbf{u}^\top \mathbf{x}^s \quad \forall \mathbf{s} \in \widetilde{\mathcal{S}}^K. \end{aligned} \quad (\text{EC.11})$$

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

$$\begin{aligned} & \text{minimize} \quad \mathbf{u}^\top \mathbf{x}^s \\ & \text{subject to} \quad \mathbf{u} \in \mathbb{R}^J, \boldsymbol{\epsilon} \in \mathbb{R}_+^K \\ & \quad \quad \quad \mathbf{u}^\top [\mathbf{s}_\kappa (\mathbf{x}_1^{\kappa_1} - \mathbf{x}_2^{\kappa_2})] + \boldsymbol{\epsilon}_\kappa \geq 0 \quad \forall \kappa \in \mathcal{K} \\ & \quad \quad \quad \mathbf{B}\mathbf{u} \geq \mathbf{b} \\ & \quad \quad \quad \mathbf{e}^\top \boldsymbol{\epsilon} \leq \Gamma. \end{aligned}$$

Its dual is expressible as

$$\begin{aligned}
& \text{maximize} && \mathbf{b}^\top \boldsymbol{\beta}^s + \Gamma \mu^s \\
& \text{subject to} && \boldsymbol{\alpha}^s \in \mathbb{R}_+^K, \boldsymbol{\beta}^s \in \mathbb{R}_+^M, \mu^s \in \mathbb{R}_- \\
& && \sum_{\kappa \in \mathcal{K}} \mathbf{s}_\kappa (\mathbf{x}^{\iota_1^\kappa} - \mathbf{x}^{\iota_2^\kappa}) \boldsymbol{\alpha}_\kappa^s + \mathbf{B}^\top \boldsymbol{\beta}^s = \mathbf{x}^s \\
& && \boldsymbol{\alpha}^s + \mu^s \mathbf{e} \leq \mathbf{0}.
\end{aligned}$$

Next, we claim that the primal-dual pair above satisfies strong duality. If $\mathbf{s} \in \mathcal{S}_\Gamma(\boldsymbol{\iota})$, then $\mathcal{U}_\Gamma(\boldsymbol{\iota}, \mathbf{s})$ and thus also $\tilde{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \mathbf{s})$ are non-empty. Since $\tilde{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \mathbf{s})$ is non-empty and compact, by Assumption 1, it follows that the primal problem above is feasible and bounded, so that it is solvable. We show that the statement also holds if $\mathbf{s} \notin \mathcal{S}_\Gamma(\boldsymbol{\iota})$. If $\mathbf{s} \notin \mathcal{S}_\Gamma(\boldsymbol{\iota})$, then the primal problem is infeasible. This implies that the dual is either infeasible or unbounded. We show that it cannot be infeasible. Fix $\boldsymbol{\alpha} = \mathbf{0}$ and $\mu = 0$ in the dual. Then, the dual reduces to

$$\begin{aligned}
& \text{maximize} && \mathbf{b}^\top \boldsymbol{\beta} \\
& \text{subject to} && \boldsymbol{\beta} \in \mathbb{R}_+^M \\
& && \mathbf{B}^\top \boldsymbol{\beta} = \mathbf{x}.
\end{aligned}$$

By Farkas' lemma, exactly one of the following alternatives must hold: *a*) There exists $\boldsymbol{\beta} \in \mathbb{R}_+^M$ such that $\mathbf{B}^\top \boldsymbol{\beta} = \mathbf{x}$; or *b*) There exists $\mathbf{u} \in \mathbb{R}^J$ such that $\mathbf{B}\mathbf{u} \geq \mathbf{0}$ and $\mathbf{u}^\top \mathbf{x} < 0$. Since the set \mathcal{U}^0 is bounded, by Assumption 1, 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 $\mathbf{s} \in \tilde{\mathcal{S}}^K$.

Replacing each minimization subproblem in Problem (EC.11) with its dual, we obtain the equivalent reformulation

$$\left. \begin{aligned}
& \text{maximize} && \tau \\
& \text{subject to} && \tau \in \mathbb{R}, \boldsymbol{\iota} \in \mathcal{C}^K \\
& && \mathbf{x}^s \in \mathcal{R}, \boldsymbol{\alpha}^s \in \mathbb{R}_+^K, \boldsymbol{\beta}^s \in \mathbb{R}_+^M, \mu^s \in \mathbb{R}_- \\
& && \tau \leq \mathbf{b}^\top \boldsymbol{\beta}^s + \Gamma \mu^s \\
& && \sum_{\kappa \in \mathcal{K}} \mathbf{s}_\kappa (\mathbf{x}^{\iota_1^\kappa} - \mathbf{x}^{\iota_2^\kappa}) \boldsymbol{\alpha}_\kappa^s + \mathbf{B}^\top \boldsymbol{\beta}^s = \mathbf{x}^s \\
& && \boldsymbol{\alpha}^s + \mu^s \mathbf{e} \leq \mathbf{0}
\end{aligned} \right\} \forall \mathbf{s} \in \tilde{\mathcal{S}}^K,$$

which concludes the proof. \square

Proof of Theorem 2 We begin by showing that Problem (1) is equivalent to

$$\begin{aligned}
& \text{maximize} && \tau \\
& \text{subject to} && \tau \in \mathbb{R}, \mathbf{v}^\kappa, \mathbf{w}^\kappa \in \{0, 1\}^I, \kappa \in \mathcal{K} \\
& && \boldsymbol{\alpha}^s \in \mathbb{R}_+^K, \boldsymbol{\beta}^s \in \mathbb{R}_+^M, \mu^s \in \mathbb{R}_-, \mathbf{x}^s \in \mathcal{R}, \mathbf{s} \in \tilde{\mathcal{S}}^K \\
& && \left. \begin{aligned} \tau &\leq \mathbf{b}^\top \boldsymbol{\beta}^s + \Gamma \mu^s \\ \sum_{\kappa \in \mathcal{K}} \mathbf{s}_\kappa \sum_{i \in \mathcal{I}} \mathbf{x}^i (\mathbf{v}_i^\kappa - \mathbf{w}_i^\kappa) \boldsymbol{\alpha}_\kappa^s + \mathbf{B}^\top \boldsymbol{\beta}^s &= \mathbf{x}^s \\ \boldsymbol{\alpha}^s + \mu^s \mathbf{e} &\leq \mathbf{0} \end{aligned} \right\} \forall \mathbf{s} \in \tilde{\mathcal{S}}^K \quad (\text{EC.12}) \\
& && \left. \begin{aligned} \mathbf{e}^\top \mathbf{v}^\kappa = 1, \mathbf{e}^\top \mathbf{w}^\kappa &= 1 \\ 1 - \mathbf{w}_i^\kappa &\geq \sum_{i': i' \geq i} \mathbf{v}_{i'}^\kappa \quad \forall i \in \mathcal{I} \end{aligned} \right\} \forall \kappa \in \mathcal{K}.
\end{aligned}$$

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

$$\mathbf{v}_i^\kappa := \begin{cases} 1 & \text{if } \boldsymbol{\nu}_1^\kappa = i \\ 0 & \text{else,} \end{cases} \quad \text{and} \quad \mathbf{w}_i^\kappa := \begin{cases} 1 & \text{if } \boldsymbol{\nu}_2^\kappa = i \\ 0 & \text{else,} \end{cases}$$

for each $i \in \mathcal{I}$. Fix $\kappa \in \mathcal{K}$. Since $\boldsymbol{\nu}^\kappa \in \mathcal{C}$, it readily follows that $\boldsymbol{\nu}_1^\kappa$ and $\boldsymbol{\nu}_2^\kappa$ are both in \mathcal{I} , implying that $\mathbf{e}^\top \mathbf{v}^\kappa = \mathbf{e}^\top \mathbf{w}^\kappa = 1$. Fix $i \in \mathcal{I}$. If $\mathbf{w}_i^\kappa = 0$, then it holds that $1 - \mathbf{w}_i^\kappa \geq \sum_{i': i' \geq i} \mathbf{v}_{i'}^\kappa$. On the other hand, if $\mathbf{w}_i^\kappa = 1$, then it holds that $\boldsymbol{\nu}_2^\kappa = i$. But since $\boldsymbol{\nu}^\kappa \in \mathcal{C}$, it holds that $\boldsymbol{\nu}_1^\kappa < \boldsymbol{\nu}_2^\kappa = i$, i.e., there must exist $i' < i$ such that $\boldsymbol{\nu}_1^\kappa = i'$ and $\mathbf{v}_{i'}^\kappa = 1$. This in turn implies that $\sum_{i': i' \geq i} \mathbf{v}_{i'}^\kappa = 0$, and therefore $1 - \mathbf{w}_i^\kappa \geq \sum_{i': i' \geq i} \mathbf{v}_{i'}^\kappa$ holds. Since the choice of $\kappa \in \mathcal{K}$ and $i \in \mathcal{I}$ was arbitrary, it holds that

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

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

$$\sum_{i \in \mathcal{I}} \mathbf{x}^i (\mathbf{v}_i^\kappa - \mathbf{w}_i^\kappa) \boldsymbol{\alpha}_\kappa^s = (\mathbf{x}^{\boldsymbol{\nu}_1^\kappa} - \mathbf{x}^{\boldsymbol{\nu}_2^\kappa}) \boldsymbol{\alpha}_\kappa^s$$

holds trivially by definition of \mathbf{v}^κ and \mathbf{w}^κ . Thus, $(\tau, \{\mathbf{v}^\kappa, \mathbf{w}^\kappa\}_{\kappa \in \mathcal{K}}, \{\boldsymbol{\alpha}^s, \boldsymbol{\beta}^s, \mu^s, \mathbf{x}^s\}_{s \in \tilde{\mathcal{S}}^K})$ is feasible in Problem (EC.12) with objective value equal to τ .

For the other direction, let $(\tau, \{\mathbf{v}^\kappa, \mathbf{w}^\kappa\}_{\kappa \in \mathcal{K}}, \{\boldsymbol{\alpha}^s, \boldsymbol{\beta}^s, \mu^s, \mathbf{x}^s\}_{s \in \tilde{\mathcal{S}}^K})$ be feasible in Problem (EC.12). Define

$$\boldsymbol{\nu}_1^\kappa = \sum_{i \in \mathcal{I}} i \cdot \mathbb{I}(\mathbf{v}_i^\kappa = 1) \quad \text{and} \quad \boldsymbol{\nu}_2^\kappa = \sum_{i \in \mathcal{I}} i \cdot \mathbb{I}(\mathbf{w}_i^\kappa = 1).$$

Fix $\kappa \in \mathcal{K}$. Then, $\boldsymbol{\nu}_1^\kappa$ and $\boldsymbol{\nu}_2^\kappa$ are both in the set \mathcal{I} . It follows from

$$1 - \mathbf{w}_i^\kappa \geq \sum_{i': i' \geq i} \mathbf{v}_{i'}^\kappa \quad \forall i \in \mathcal{I}$$

that if $\mathbf{w}_i^\kappa = 1$, then $\iota_2^\kappa = i$ and $\sum_{i':i' \geq i} \mathbf{v}_{i'}^\kappa = 0$, implying that $\sum_{i':i' < i} \mathbf{v}_{i'}^\kappa = 1$, i.e., $\iota_1^\kappa < i = \iota_2^\kappa$ and it holds that $\iota^\kappa \in \mathcal{C}$. Since the choice of κ was arbitrary, $\iota \in \mathcal{C}^K$. Fix $\mathbf{s} \in \tilde{\mathcal{S}}^K$ and $\kappa \in \mathcal{K}$. Then,

$$(\mathbf{x}^{\iota_1^\kappa} - \mathbf{x}^{\iota_2^\kappa}) \boldsymbol{\alpha}_\kappa^{\mathbf{s}} = \sum_{i \in \mathcal{I}} \mathbf{x}^i (\mathbf{v}_i^\kappa - \mathbf{w}_i^\kappa) \boldsymbol{\alpha}_\kappa^{\mathbf{s}}$$

holds trivially by definition of ι^κ . Therefore, $(\tau, \iota, \{\boldsymbol{\alpha}^{\mathbf{s}}, \boldsymbol{\beta}^{\mathbf{s}}, \boldsymbol{\mu}^{\mathbf{s}}, \mathbf{x}^{\mathbf{s}}\}_{\mathbf{s} \in \tilde{\mathcal{S}}^K})$ is feasible in Problem (1) with objective value equal to τ . We have thus shown that Problems (1) and (EC.12) are equivalent.

Equivalence of the bilinear problem (EC.12) and the MILP (2) 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 1 Since ι is feasible in Problem $(\mathcal{CCG}_u^{\text{main}}(\mathcal{S}'))$, it follows that $\iota \in \mathcal{C}^K$. By following a proof strategy similar to that in the Proof of Lemma 2 (based on Farkas' lemma), it follows that ι is feasible in Problem $(\widetilde{\mathcal{QP}}_u^K)$. Thus, it remains to show that for any given $\iota \in \mathcal{C}^K$, Problem $(\mathcal{CCG}_u^{\text{feas}}(\iota))$ is equivalent to

$$\text{minimize}_{\mathbf{s} \in \tilde{\mathcal{S}}^K} \max_{\mathbf{x} \in \mathcal{R}} \min_{\mathbf{u} \in \tilde{\mathcal{U}}_\Gamma(\iota, \mathbf{s})} \mathbf{u}^\top \mathbf{x} \quad (\text{EC.13})$$

in the sense that the two problems have the same optimal objective value. To this end, fix $\iota \in \mathcal{C}^K$. Using an epigraph reformulation, we can write Problem (EC.13) equivalently as

$$\begin{aligned} & \text{minimize} && \theta \\ & \text{subject to} && \theta \in \mathbb{R}, \mathbf{s} \in \tilde{\mathcal{S}}^K \\ & && \theta \geq \max_{\mathbf{x} \in \mathcal{R}} \min_{\mathbf{u} \in \tilde{\mathcal{U}}_\Gamma(\iota, \mathbf{s})} \mathbf{u}^\top \mathbf{x}. \end{aligned} \quad (\text{EC.14})$$

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

$$\begin{aligned} & \text{minimize} && \theta \\ & \text{subject to} && \theta \in \mathbb{R}, \mathbf{s} \in \tilde{\mathcal{S}}^K \\ & && \mathbf{u}^{\mathbf{x}} \in \tilde{\mathcal{U}}_\Gamma(\iota, \mathbf{s}) \quad \forall \mathbf{x} \in \mathcal{R} \\ & && \theta \geq (\mathbf{u}^{\mathbf{x}})^\top \mathbf{x} \quad \forall \mathbf{x} \in \mathcal{R}. \end{aligned}$$

From the definition of $\tilde{\mathcal{U}}_\Gamma(\iota, \mathbf{s})$, we can rewrite the preceding problem equivalently as

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

which, for M sufficiently large, is equivalent to Problem $(\mathcal{CCG}_u^{\text{feas}}(\iota))$. Thus, Problems (EC.13) and $(\mathcal{CCG}_u^{\text{feas}}(\iota))$ are equivalent. Since the choice of ι was arbitrary, the claim follows. \square

Proof of Lemma 3 (i) By virtue of Proposition 1, θ yields a lower bound to the optimal value of Problem $(\widetilde{\mathcal{QP}}_u^K)$. 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 $(\widetilde{\mathcal{QP}}_u^K)$. Therefore, $\theta \leq \tau$.

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

$$\begin{aligned} & \text{maximize} && \mathbf{b}^\top \boldsymbol{\beta} + \Gamma \mu \\ & \text{subject to} && \mathbf{x} \in \mathcal{R}, \boldsymbol{\alpha} \in \mathbb{R}^K, \boldsymbol{\beta} \in \mathbb{R}_+^M, \mu \in \mathbb{R}_- \\ & && \sum_{\kappa \in \mathcal{K}} \mathbf{s}_\kappa (\mathbf{x}^{\iota_1^\kappa} - \mathbf{x}^{\iota_2^\kappa}) \boldsymbol{\alpha}_\kappa + \mathbf{B}^\top \boldsymbol{\beta} = \mathbf{x} \\ & && \boldsymbol{\alpha} + \mu \mathbf{e} \leq \mathbf{0}. \end{aligned} \tag{EC.15}$$

An inspection of the proof of Proposition 1 reveals that Problems $(\mathcal{CCG}_u^{\text{feas}}(\boldsymbol{\iota}))$ and (EC.13) are equivalent. An inspection of the proof of Lemma 2 shows that Problem (EC.13) is equivalent to

$$\begin{aligned} & \text{minimize} && \max_{\mathbf{s} \in \widetilde{\mathcal{S}}^K} \mathbf{b}^\top \boldsymbol{\beta} + \Gamma \mu \\ & && \text{s. t.} \quad \mathbf{x} \in \mathcal{R}, \boldsymbol{\alpha} \in \mathbb{R}^K, \boldsymbol{\beta} \in \mathbb{R}_+^M, \mu \in \mathbb{R}_- \\ & && \sum_{\kappa \in \mathcal{K}} \mathbf{s}_\kappa (\mathbf{x}^{\iota_1^\kappa} - \mathbf{x}^{\iota_2^\kappa}) \boldsymbol{\alpha}_\kappa + \mathbf{B}^\top \boldsymbol{\beta} = \mathbf{x} \\ & && \boldsymbol{\alpha} + \mu \mathbf{e} \leq \mathbf{0}. \end{aligned} \tag{EC.16}$$

Thus, Problems (EC.16) and $(\mathcal{CCG}_u^{\text{feas}}(\boldsymbol{\iota}))$ are equivalent and Problem (EC.16) has an optimal objective value of θ . This implies that $\theta < \tau$, a contradiction.

(iii) Suppose $\theta < \tau$, and let \mathbf{s} be defined as in the premise of the lemma. Then, the proof of item (ii) reveals that \mathbf{s} must be optimal in Problem (EC.16) with associated optimal value θ . This in turn implies that \mathbf{s} is such that the optimal objective value of Problem (EC.15) is strictly less than τ , implying that the \mathbf{s} th subproblem $(\mathcal{CCG}_u^{\text{sub},\mathbf{s}}(\tau, \boldsymbol{\iota}))$ 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 $\text{UB} - \text{LB} \leq \delta$ (in which case the algorithm terminates) or a new set of constraints (indexed by the infeasible index \mathbf{s}) is added to the master problem $\mathcal{CCG}_u^{\text{main}}(\mathcal{S}')$, 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 $\text{UB} - \text{LB} \leq \delta$, we are guaranteed that, at termination, the returned solution will have an objective value that is within δ of the optimal objective value of the problem. This concludes the proof. \square

EC.8. Proofs of Statements in Sections 5

The following lemma shows that, in a manner paralleling the risk averse case, Problem (QP_r^K) 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 $\mathcal{U}(\boldsymbol{\iota}, \mathbf{s})$ by their loose counterparts. It is needed in the proof of Theorem 4.

Lemma EC.1. *Problem (QP_r^K) is equivalent to*

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

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

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

Proof of Theorem 4 We work with the equivalent reformulation $(\tilde{\mathcal{P}}_r^K)$ of (QP_r^K) , see Lemma EC.1 and 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 $(\tilde{\mathcal{P}}_r^K)$. For convenience, we work with the negative of the objective function of Problem $(\tilde{\mathcal{P}}_r^K)$ given by

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

To this end, fix an instance (n, \mathbf{w}, W) of the partition problem. Set $\Gamma = 0$, $J = 2n + 3$, $K = n$, and

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

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 $\mathbf{s} \in \tilde{\mathcal{S}}^K$, we have

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

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

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

Then, we have

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

Next, observe that

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

Similarly, it holds that

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

Therefore,

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

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 $(\widetilde{\mathcal{QP}}_u^K)$ is $-4W$. To this end, note that the objective value of $\boldsymbol{\iota}$ is given by

$$\underset{\mathbf{s} \in \tilde{\mathcal{S}}^K}{\text{minimize}} \quad Z(\mathbf{s}) \tag{EC.17}$$

and is lower bounded by $-4W$. If there exists $\mathcal{X} \subset \mathcal{A}$ such that $\sum_{i \in \mathcal{X}} \mathbf{w}_i = W$, then the solution $\mathbf{s} \in \tilde{\mathcal{S}}^K$ defined through $s_\kappa = 1$ if $\kappa \in \mathcal{X}$, $= -1$ else, $\kappa \in \mathcal{K}$, attains an objective value of $-4W$ in Problem (EC.17). Conversely, if the optimal objective value of Problem (EC.17) is $-4W$, then the set $\mathcal{X} = \{\kappa \in \mathcal{A} : s_\kappa = 1\}$ is such that $\sum_{i \in \mathcal{X}} \mathbf{w}_i = W$ and the claim follows. This concludes the proof of the second item. \square

The following Lemma is needed in the proof of Theorem 5.

Lemma EC.2. *Problem $(\tilde{\mathcal{P}}_r^K)$ is equivalent to the following finite program*

$$\begin{array}{ll}
\text{minimize} & \tau \\
\text{subject to} & \tau \in \mathbb{R}, \boldsymbol{\nu} \in \mathcal{C}^K, \mathbf{x}^s \in \mathcal{R}, \mathbf{s} \in \tilde{\mathcal{S}}^K \\
& \left. \begin{array}{l}
\boldsymbol{\alpha}^{(\mathbf{x}', \mathbf{s})} \in \mathbb{R}_-^K, \boldsymbol{\beta}^{(\mathbf{x}', \mathbf{s})} \in \mathbb{R}_-^M, \mu^{(\mathbf{x}', \mathbf{s})} \in \mathbb{R}_+ \\
\tau \geq \mathbf{b}^\top \boldsymbol{\beta}^{(\mathbf{x}', \mathbf{s})} + \Gamma \mu^{(\mathbf{x}', \mathbf{s})} \\
\sum_{\kappa \in \mathcal{K}} s_\kappa (\mathbf{x}^{\kappa_1} - \mathbf{x}^{\kappa_2}) \boldsymbol{\alpha}_\kappa^{(\mathbf{x}', \mathbf{s})} + \mathbf{B}^\top \boldsymbol{\beta}^{(\mathbf{x}', \mathbf{s})} = \mathbf{x}' - \mathbf{x}^s \\
\boldsymbol{\alpha}^{(\mathbf{x}', \mathbf{s})} + \mu^{(\mathbf{x}', \mathbf{s})} \mathbf{e} \geq \mathbf{0}
\end{array} \right\} \quad \forall \mathbf{s} \in \tilde{\mathcal{S}}^K, \forall \mathbf{x}' \in \mathcal{R},
\end{array} \tag{EC.18}$$

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

Proof of Lemma EC.2 In a way that parallels the proof of Observation EC.1, it can be readily shown that Problem $(\tilde{\mathcal{P}}_r^K)$ is equivalent to

$$\text{minimize}_{\boldsymbol{\nu} \in \mathcal{C}^K} \min_{\substack{\mathbf{x}^s \in \mathcal{R}: \\ \mathbf{s} \in \tilde{\mathcal{S}}^K}} \max_{\mathbf{s} \in \tilde{\mathcal{S}}^K} \max_{\mathbf{u} \in \tilde{\mathcal{U}}_\Gamma(\boldsymbol{\nu}, \mathbf{s})} \left\{ \max_{\mathbf{x}' \in \mathcal{R}} \mathbf{u}^\top \mathbf{x}' - \mathbf{u}^\top \mathbf{x}^s \right\}.$$

The above problem can be written in epigraph form equivalently as

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

which is in turn equivalent to

$$\begin{array}{ll}
\text{minimize} & \tau \\
\text{subject to} & \tau \in \mathbb{R}, \boldsymbol{\nu} \in \mathcal{C}^K, \mathbf{x}^s \in \mathcal{R}, \mathbf{s} \in \tilde{\mathcal{S}}^K \\
& \tau \geq \max_{\mathbf{u} \in \tilde{\mathcal{U}}_\Gamma(\boldsymbol{\nu}, \mathbf{s})} \mathbf{u}^\top (\mathbf{x}' - \mathbf{x}^s) \quad \forall \mathbf{s} \in \tilde{\mathcal{S}}^K, \forall \mathbf{x}' \in \mathcal{R}.
\end{array} \tag{EC.19}$$

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

$$\begin{aligned} & \text{maximize} && \mathbf{u}^\top (\mathbf{x}' - \mathbf{x}^s) \\ & \text{subject to} && \mathbf{u} \in \mathbb{R}^J, \boldsymbol{\epsilon} \in \mathbb{R}_+^K \\ & && \mathbf{u}^\top [\mathbf{s}_\kappa (\mathbf{x}'^{\iota_\kappa^1} - \mathbf{x}'^{\iota_\kappa^2})] + \boldsymbol{\epsilon}_\kappa \geq 0 \quad \forall \kappa \in \mathcal{K} \\ & && \mathbf{B}\mathbf{u} \geq \mathbf{b} \\ & && \mathbf{e}^\top \boldsymbol{\epsilon} \leq \Gamma. \end{aligned}$$

Its dual is expressible as

$$\begin{aligned} & \text{minimize} && \mathbf{b}^\top \boldsymbol{\beta}^{(\mathbf{x}', \mathbf{s})} + \Gamma \mu^{(\mathbf{x}', \mathbf{s})} \\ & \text{subject to} && \boldsymbol{\alpha}^{(\mathbf{x}', \mathbf{s})} \in \mathbb{R}_-^K, \boldsymbol{\beta}^{(\mathbf{x}', \mathbf{s})} \in \mathbb{R}_-^M, \mu^{(\mathbf{x}', \mathbf{s})} \in \mathbb{R}_+ \\ & && \sum_{\kappa \in \mathcal{K}} \mathbf{s}_\kappa (\mathbf{x}'^{\iota_\kappa^1} - \mathbf{x}'^{\iota_\kappa^2}) \boldsymbol{\alpha}_\kappa^{(\mathbf{x}', \mathbf{s})} + \mathbf{B}^\top \boldsymbol{\beta}^s = \mathbf{x}' - \mathbf{x}^s \\ & && \boldsymbol{\alpha}^{(\mathbf{x}', \mathbf{s})} + \mu^{(\mathbf{x}', \mathbf{s})} \mathbf{e} \geq \mathbf{0}. \end{aligned}$$

Following a proof strategy similar to that taken in the proof of Lemma 2, 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.19) with its dual yields the equivalent formulation

$$\begin{aligned} & \text{minimize} && \tau \\ & \text{subject to} && \tau \in \mathbb{R}, \boldsymbol{\iota} \in \mathcal{C}^K, \mathbf{x}^s \in \mathcal{R}, \mathbf{s} \in \tilde{\mathcal{S}}^K \\ & && \boldsymbol{\alpha}^{(\mathbf{x}', \mathbf{s})} \in \mathbb{R}_-^K, \boldsymbol{\beta}^{(\mathbf{x}', \mathbf{s})} \in \mathbb{R}_-^M, \mu^{(\mathbf{x}', \mathbf{s})} \in \mathbb{R}_+ \\ & && \tau \geq \mathbf{b}^\top \boldsymbol{\beta}^{(\mathbf{x}', \mathbf{s})} + \Gamma \mu^{(\mathbf{x}', \mathbf{s})} \\ & && \sum_{\kappa \in \mathcal{K}} \mathbf{s}_\kappa (\mathbf{x}'^{\iota_\kappa^1} - \mathbf{x}'^{\iota_\kappa^2}) \boldsymbol{\alpha}_\kappa^{(\mathbf{x}', \mathbf{s})} + \mathbf{B}^\top \boldsymbol{\beta}^{(\mathbf{x}', \mathbf{s})} = \mathbf{x}' - \mathbf{x}^s \\ & && \boldsymbol{\alpha}^{(\mathbf{x}', \mathbf{s})} + \mu^{(\mathbf{x}', \mathbf{s})} \mathbf{e} \geq \mathbf{0} \end{aligned} \quad \left. \vphantom{\begin{aligned} & \text{subject to} \\ & \boldsymbol{\alpha}^{(\mathbf{x}', \mathbf{s})} \in \mathbb{R}_-^K, \boldsymbol{\beta}^{(\mathbf{x}', \mathbf{s})} \in \mathbb{R}_-^M, \mu^{(\mathbf{x}', \mathbf{s})} \in \mathbb{R}_+ \\ & \tau \geq \mathbf{b}^\top \boldsymbol{\beta}^{(\mathbf{x}', \mathbf{s})} + \Gamma \mu^{(\mathbf{x}', \mathbf{s})} \\ & \sum_{\kappa \in \mathcal{K}} \mathbf{s}_\kappa (\mathbf{x}'^{\iota_\kappa^1} - \mathbf{x}'^{\iota_\kappa^2}) \boldsymbol{\alpha}_\kappa^{(\mathbf{x}', \mathbf{s})} + \mathbf{B}^\top \boldsymbol{\beta}^{(\mathbf{x}', \mathbf{s})} = \mathbf{x}' - \mathbf{x}^s \\ & \boldsymbol{\alpha}^{(\mathbf{x}', \mathbf{s})} + \mu^{(\mathbf{x}', \mathbf{s})} \mathbf{e} \geq \mathbf{0} \end{aligned}} \right\} \forall \mathbf{s} \in \tilde{\mathcal{S}}^K, \forall \mathbf{x}' \in \mathcal{R}. \quad (\text{EC.20})$$

This concludes the proof. \square

Proof of Theorem 5 The proof follows directly from Lemma EC.2 by following an approach similar to that taken in the Proof of Theorem 2 and is thus omitted. \square

We introduce the following problem which is needed in the proof of Theorem EC.1. Given variables $(\tau, \boldsymbol{\iota})$ feasible in the main problem, we define the $(\mathbf{x}', \mathbf{s})$ th subproblem, $\mathbf{x}' \in \mathcal{R}$, $\mathbf{s} \in \tilde{\mathcal{S}}^K$, through

$$\begin{aligned} & \text{minimize} && 0 \\ & \text{subject to} && \boldsymbol{\alpha} \in \mathbb{R}_-^K, \boldsymbol{\beta} \in \mathbb{R}_-^M, \mu \in \mathbb{R}_+, \mathbf{x} \in \mathcal{R} \\ & && \tau \geq \mathbf{b}^\top \boldsymbol{\beta} + \Gamma \mu \\ & && \sum_{\kappa \in \mathcal{K}} \mathbf{s}_\kappa (\mathbf{x}'^{\iota_\kappa^1} - \mathbf{x}'^{\iota_\kappa^2}) \boldsymbol{\alpha}_\kappa + \mathbf{B}^\top \boldsymbol{\beta} = \mathbf{x}' - \mathbf{x} \\ & && \boldsymbol{\alpha} + \mu \mathbf{e} \geq \mathbf{0}. \end{aligned} \quad (\mathcal{CCG}_{\text{mmr}}^{\text{sub}, (\mathbf{x}', \mathbf{s})}(\tau, \boldsymbol{\iota}))$$

The following proposition and lemma are needed in the proof of Theorem EC.1.

Proposition EC.2. *Let $\boldsymbol{\iota}$ be feasible in the relaxed master problem $(\mathcal{CCG}_r^{\text{main}}(\mathcal{S}'))$. Then, $\boldsymbol{\iota}$ is feasible in Problem $(\tilde{\mathcal{P}}_r^K)$ and the objective value of $\boldsymbol{\iota}$ in Problem $(\tilde{\mathcal{P}}_r^K)$ is given by the optimal objective value of Problem $(\mathcal{CCG}_r^{\text{feas}}(\boldsymbol{\iota}))$. Moreover, if the set \mathcal{R} has fixed finite cardinality, then Problem $(\mathcal{CCG}_r^{\text{feas}}(\boldsymbol{\iota}))$ is a mixed-binary linear program of size polynomial in the size of the input.*

Proof of Proposition EC.2 Since $\boldsymbol{\iota}$ is feasible in Problem $(\mathcal{CCG}_r^{\text{main}}(\mathcal{S}'))$, it follows that $\boldsymbol{\iota} \in \mathcal{C}^K$. Thus, $\boldsymbol{\iota}$ is feasible in Problem $(\tilde{\mathcal{P}}_r^K)$. It remains to show that for any given $\boldsymbol{\iota} \in \mathcal{C}^K$, Problems $(\mathcal{CCG}_r^{\text{feas}}(\boldsymbol{\iota}))$ and

$$\max_{\boldsymbol{s} \in \tilde{\mathcal{S}}^K} \min_{\boldsymbol{x} \in \mathcal{R}} \max_{\boldsymbol{u} \in \tilde{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s})} \left\{ \max_{\boldsymbol{x}' \in \mathcal{R}} \boldsymbol{u}^\top \boldsymbol{x}' - \boldsymbol{u}^\top \boldsymbol{x} \right\} \quad (\text{EC.21})$$

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.21) equivalently as

$$\begin{aligned} & \text{maximize} && \theta \\ & \text{subject to} && \theta \in \mathbb{R}, \boldsymbol{s} \in \tilde{\mathcal{S}}^K \\ & && \theta \leq \min_{\boldsymbol{x} \in \mathcal{R}} \max_{\boldsymbol{u} \in \tilde{\mathcal{U}}_\Gamma(\boldsymbol{\iota}, \boldsymbol{s})} \max_{\boldsymbol{x}' \in \mathcal{R}} \boldsymbol{u}^\top \boldsymbol{x}' - \boldsymbol{u}^\top \boldsymbol{x}, \end{aligned} \quad (\text{EC.22})$$

and thus as

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

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

$$\begin{aligned} & \text{maximize} && \theta \\ & \text{subject to} && \theta \in \mathbb{R}, \boldsymbol{s} \in \tilde{\mathcal{S}}^K \\ & && \boldsymbol{x}'^{\boldsymbol{x}} \in \mathcal{R}, \boldsymbol{u}^{\boldsymbol{x}} \in \mathcal{U}^0, \boldsymbol{\epsilon}^{\boldsymbol{x}} \in \mathcal{E}_\Gamma \quad \forall \boldsymbol{x} \in \mathcal{R} \\ & && \theta \leq (\boldsymbol{u}^{\boldsymbol{x}})^\top (\boldsymbol{x}'^{\boldsymbol{x}} - \boldsymbol{x}) \quad \forall \boldsymbol{x} \in \mathcal{R} \\ & && \left. \begin{aligned} & (\boldsymbol{u}^{\boldsymbol{x}})^\top (\boldsymbol{x}^{\iota_k} - \boldsymbol{x}'^{\iota_k}) \geq -\boldsymbol{\epsilon}_\kappa^{\boldsymbol{x}} \quad \forall \kappa \in \mathcal{K} : \boldsymbol{s}_\kappa = 1 \\ & (\boldsymbol{u}^{\boldsymbol{x}})^\top (\boldsymbol{x}^{\iota_k} - \boldsymbol{x}'^{\iota_k}) \leq \boldsymbol{\epsilon}_\kappa^{\boldsymbol{x}} \quad \forall \kappa \in \mathcal{K} : \boldsymbol{s}_\kappa = -1 \end{aligned} \right\} \quad \forall \boldsymbol{x} \in \mathcal{R}. \end{aligned}$$

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

Lemma EC.3. *The relaxed master problem $(\mathcal{CCG}_r^{\text{main}}(\mathcal{S}'))$ is always feasible. If $(\mathcal{CCG}_r^{\text{main}}(\mathcal{S}'))$ 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 $(\mathcal{CCG}_r^{\text{main}}(\mathcal{S}'))$ is unbounded, set $\tau = -\infty$ and let $\boldsymbol{\iota} \in \mathcal{C}^K$ be such that $(\mathcal{CCG}_r^{\text{main}}(\mathcal{S}'))$ 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 $(\mathcal{CCG}_r^{\text{feas}}(\boldsymbol{\iota}))$. Then, the following hold:*

- (i) $\theta \geq \tau$;
- (ii) If $\theta = \tau$, then Problem $(\mathcal{CCG}_{\text{mmr}}^{\text{sub},(\mathbf{x}',\mathbf{s})}(\tau,\boldsymbol{\iota}))$ is feasible for all $\mathbf{x}' \in \mathcal{R}$ and $\mathbf{s} \in \tilde{\mathcal{S}}^K$;
- (iii) If $\theta > \tau$, then there exists $\mathbf{x}' \in \mathcal{R}$ such that the pair (\mathbf{x}',\mathbf{s}) corresponds to an infeasible subproblem, i.e., Problem $(\mathcal{CCG}_{\text{mmr}}^{\text{sub},(\mathbf{x}',\mathbf{s})}(\tau,\boldsymbol{\iota}))$ is infeasible.

Proof of Lemma EC.3 (i) By virtue of Proposition EC.2, θ yields an upper bound to the optimal value of the Problem $(\tilde{\mathcal{P}}_r^K)$. At the same time, since $\tilde{\mathcal{R}} \subseteq \mathcal{R}$ and $\mathcal{S}' \subseteq \tilde{\mathcal{S}}^K$, it is evident that τ is a lower bound on the optimal objective value of Problem $(\tilde{\mathcal{P}}_r^K)$. Therefore, $\theta \geq \tau$.

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

$$\begin{aligned}
& \text{minimize} && \mathbf{b}^\top \boldsymbol{\beta} + \Gamma \mu \\
& \text{subject to} && \boldsymbol{\alpha} \in \mathbb{R}_-^K, \boldsymbol{\beta} \in \mathbb{R}_-^M, \mu \in \mathbb{R}_+, \mathbf{x} \in \mathcal{R} \\
& && \sum_{\kappa \in \mathcal{K}} \mathbf{s}_\kappa (\mathbf{x}^{\iota_1^\kappa} - \mathbf{x}^{\iota_2^\kappa}) \boldsymbol{\alpha}_\kappa + \mathbf{B}^\top \boldsymbol{\beta} = \mathbf{x}' - \mathbf{x} \\
& && \boldsymbol{\alpha} + \mu \mathbf{e} \geq \mathbf{0}.
\end{aligned} \tag{EC.23}$$

An inspection of the proof of Proposition EC.2 reveals that Problems $(\mathcal{CCG}_r^{\text{feas}}(\boldsymbol{\iota}))$ and (EC.21) are equivalent. An inspection of the proof of Lemma EC.2 shows that Problem (EC.21) is equivalent to

$$\begin{aligned}
& \text{maximize}_{\mathbf{x}' \in \mathcal{R}, \mathbf{s} \in \tilde{\mathcal{S}}^K} && \min && \mathbf{b}^\top \boldsymbol{\beta} + \Gamma \mu \\
& \text{s. t.} && && \boldsymbol{\alpha} \in \mathbb{R}_-^K, \boldsymbol{\beta} \in \mathbb{R}_-^M, \mu \in \mathbb{R}_+, \mathbf{x} \in \mathcal{R} \\
& && && \sum_{\kappa \in \mathcal{K}} \mathbf{s}_\kappa (\mathbf{x}^{\iota_1^\kappa} - \mathbf{x}^{\iota_2^\kappa}) \boldsymbol{\alpha}_\kappa + \mathbf{B}^\top \boldsymbol{\beta} = \mathbf{x}' - \mathbf{x} \\
& && && \boldsymbol{\alpha} + \mu \mathbf{e} \geq \mathbf{0}.
\end{aligned} \tag{EC.24}$$

Thus, Problems $(\mathcal{CCG}_r^{\text{feas}}(\boldsymbol{\iota}))$ and (EC.24) are equivalent and Problem (EC.24) has an optimal value of θ . This implies that $\tau < \theta$, a contradiction.

- (iii) Suppose $\theta > \tau$ and let (\mathbf{x}',\mathbf{s}) be defined as in the premise of the lemma. Then, the proof of item (ii) reveals that \mathbf{s} is optimal in (EC.21) with associated optimal value θ . This implies that \mathbf{s} is such that the optimal objective value of Problem (EC.23) is strictly greater than τ , so that the (\mathbf{x}',\mathbf{s}) th subproblem $(\mathcal{CCG}_{\text{mmr}}^{\text{sub},(\mathbf{x}',\mathbf{s})}(\tau,\boldsymbol{\iota}))$ is infeasible, which concludes the proof.

All claims are thus proved. □

Proof of Theorem EC.1 The proof mirrors the proof of Theorem 3 and is thus omitted in the interest of space. □