Fair and Welfare-Efficient Resource Allocation under Uncertainty

Cyrus Cousins, Elita Lobo, Justin Payan,[*] and Yair Zick[†]
University of Massachusetts Amherst
{ccousins, elobo, jpayan, yzick}@umass.edu

Abstract

We study fair allocation of constrained resources, where a market designer optimizes overall welfare while maintaining group fairness. In many large-scale settings, utilities are not known in advance, but are instead observed after realizing the allocation. We therefore *estimate* agent utilities using machine learning. Optimizing over estimates requires trading-off between mean utilities and their predictive variances. We discuss these trade-offs under two paradigms for preference modeling – in the *stochastic optimization* regime, the market designer has access to a probability distribution over utilities, and in the *robust optimization* regime they have access to an *uncertainty set* containing the true utilities with high probability. We discuss utilitarian and egalitarian based objectives, and we explore how to optimize for them under stochastic and robust paradigms. We demonstrate the efficacy of our approaches on three publicly available conference reviewer assignment datasets. The approaches presented enable scalable constrained resource allocation under uncertainty for many different objectives and preference models.

1 Introduction

Constrained resource allocation without money underpins many important systems; the list of applications includes reviewer assignment for peer review (our primary example throughout the paper) [4, 14, 30, 45, 54], assigning resources to homeless populations [5, 33, 49], distributing emergency response resources [51, 56, 57], and more [1, 43, 53]. In these settings we assign *resources* to *agents*. Agents and resources are *constrained*; each agent has bounds on the minimum or maximum number of items they receive from different categories, and each item has required minimums and limited total capacity. In the case of reviewer assignment, for example, papers must receive a certain number of reviews from unique reviewers, reviewers have upper limits on the number of papers they can review, and conflicts of interest prevent specific reviewers from being assigned to a paper.

A crucial factor in all of the above settings is the presence of *uncertainty*. Uncertainty often stems from the fact that agents' utilities for resources depend on future outcomes. In reviewer assignment, a reviewer's true quality is observed only after he or she has written a review. Uncertainty may also stem from our limited ability to elicit preferences; for example, in deciding where to target lead pipe mitigation projects based on number of school-aged children per neighborhood, we may have access to imperfect school enrollment records, allowing only an approximate model of the impacts of mitigation on children in each neighborhood [53]. We adopt two possible stances towards uncertainty, depending on the information available. When we have access to *a probability distribution* over preferences, we optimize the conditional expectation of the distribution at percentiles of interest [32, 50]. When we have access to *a set* of possible preferences, we adopt the *robust* approach; which is related to the minimax regret objective used in solving robust assignment problems

[*]Corresponding author.
[†]Authors ordered alphabetically by last name.

Preprint. Under review.

[3, 9, 10, 31]. Uncertainty-aware optimization approaches can often result in significantly different allocations from the default of optimizing for welfare over a central estimate (see Example 2.1 for an intuitive example).

Typically, we maximize the sum of agent utilities. However, in many of these settings, we are also concerned with *fairness* to individuals or groups of agents. Groups of agents may represent subject areas of papers in reviewer assignment, demographic groups in poverty alleviation campaigns, or regional groupings of computational resources in bandwidth allocation. Fairness to these groups may be legally required in some cases; in others it is an ethical choice by the decision maker. Although groups are often first-class objects worthy of receiving fair treatment, group fairness is often the smallest granularity of fairness achievable under uncertainty – in a large dataset uncertainty will always cause some individuals to have vanishing welfare, but group welfare can still be upheld. Although there is much literature on combinatorial optimization under uncertainty [3, 9, 10, 31, 32], to our knowledge it has not addressed the intersection of fairness and uncertainty in the constrained multi-matching problem.

1.1 Our Contributions

We study the broad problem of fair and efficient constrained multi-matchings under preference uncertainty. We present and optimize for welfare while simultaneously accounting for the uncertainty inherent in real-world resource allocation problems. Specifically, we develop methods to efficiently optimize the utilitarian and egalitarian objective using the robust approach [6, 7, 25] and CVaR approach [50]. Our results are summarized in Table 1.

For robust optimization, we construct an uncertainty set containing the true preferences with high probability (Section 3). This model is appropriate when building a predictor with statistical error bounds, but without making any assumptions on the full probability distribution over preferences. For utilitarian and egalitarian welfare functions, we robustly maximize welfare over such uncertainty sets. When the uncertainty sets are linear we can efficiently compute the exact optimal allocations for both utilitarian and egalitarian welfare in polynomial time (Propositions 3.2 and 3.5). Under ellipsoidal uncertainty sets, we can apply an *iterated quadratic programming* approach for utilitarian welfare (Proposition 3.3), while a sub-gradient ascent approach is needed for the egalitarian objective (Proposition 3.4).

When the market designer can construct a full probability distribution over preferences, we consider *stochastic optimization* using the robustness concept of *Conditional Value at Risk*, or CVaR [50]. This approach selects an allocation that maximizes the conditional expectation of welfare over the left tail of the distribution. We largely approach CVaR objectives using sampling, then solving the resulting linear program (LP). Section 4 deals with CVaR of welfare.

In addition to discussing the theoretical underpinnings of these optimization problems, we compare them empirically in Section 5 on reviewer assignment data from AAMAS 2015, 2016, and 2021.

1.2 Related work

We discuss the history of prior work on robust and CVaR optimization in Appendix A.

Some existing work applies stochastic or robust optimization to fair division problems. A line of work studies the minimax regret objective in combinatorial optimization problems, such as constrained resource allocation [3, 9, 10, 31]. This work does not explicitly consider multi-matching problems like those considered here, nor does it address the robust egalitarian welfare problem. Pujol *et al.* [48] study fair division problems with parameters noised for differential privacy, showing that the noise can cause unfair allocations; they propose a Monte Carlo approach to mitigate the unfairness with high probability. Peters *et al.* [46] study envy-free rent division under probabilistic uncertainty. Here, a central mechanism divides rooms and sets room prices for the items to minimize envy. We study a setting without money, both utilitarian and egalitarian objectives, and robust optimization in addition to stochastic optimization. Cousins *et al.* [14] introduce the framework of Robust Reviewer Assignment to solve reviewer assignment problems, but only study robust optimization under the utilitarian objective. However, they propose a naive projected sub-gradient ascent method which requires solving a quadratic program over a large number of iterations, making it inefficient. Our empirical analysis in Section 5 demonstrates the inefficiency of this method.

Fair machine learning algorithms [15, 21, 22, 42, 59] often employ similar adversarial optimization techniques over an uncertainty set in a machine learning context. Other fair allocation research has studied the case where agent demand or item availability are uncertain but preferences are known [2, 12, 20, 26] In our case demand and availability are known but preferences are not. Devic *et al.* [18] consider fair two-sided matching where the fairness constraint is defined with respect to unknown parameters; we assume knowledge of the parameters that define the fairness constraint (i.e., group identities).

2 Fair Resource Allocation under Uncertainty

Table 1: Summary of optimization algorithms for efficiently computing utilitarian and egalitarian welfare under different robustness concepts.

	Robustness Concept			
	Robust-Linear	**Robust-Ellipsoid**	**CVaR**	**CVaR-Normal**
Utilitarian	Reduction to LP (3.2)	IQP (3.3)	Sampling + LP (4.1)	Projected Gradient Ascent (4.4)
Egalitarian	Reduction to LP (3.5)	Sub-Gradient Ascent (3.4)	Sampling + LP (H.3)	SOCP (C.1)

2.1 Fair Resource Allocation

We have a set of n agents $N = \{a_1, \dots a_n\}$, and m item types $I = \{i_1, \dots i_m\}$. Agents are partitioned into g groups $\mathcal{G} = \{G_1, \dots G_g\}$, with each $G \subseteq N$ and each agent i belonging to a unique group.

Here onwards, for any given $n \times m$ matrix $\boldsymbol{X}$ we use the same lower-case bold letter, i.e., $\boldsymbol{x}$ to denote the vector representing the vectorized form of the matrix $\boldsymbol{X}$, in row-major order. Furthermore, for any group of agents G, we use $\mathbf{x}_G \in \mathbb{R}^{|G|m}$ to denote the vector restricted to the agents in G. Given vectors $\mathbf{x}, \mathbf{y} \subseteq \mathbb{R}^{nm}$ and real number $c \in \mathbb{R}$, let $\mathbf{x} \succeq c$ denote that $\mathbf{x}_j \geq c$ for all j, and let $\mathbf{x} \succeq \mathbf{y}$ denote that $\mathbf{x} - \mathbf{y} \succeq 0$. The $\preceq$ operator is defined analogously.

We assume a valuation matrix $\boldsymbol{V}^* \in [0,1]^{n \times m}$, where $\boldsymbol{V}^*_{a,i}$ encodes the true value of assigning item type i to agent a. The values of $\boldsymbol{V}^*$ are typically unknown; we discuss our approaches to handle this problem in Section 2.2.

Given some set of feasible assignments $\mathcal{A} \subseteq \mathbb{N}^{n \times m}$, we aim to find assignments $\boldsymbol{A} \in \mathcal{A}$ where $\boldsymbol{A}_{a,i}$ indicates the number of items of type i allocated to agent a. For each agent $a \in N$, we have upper and lower bounds on assignments of the form $\underline{\kappa}_a \leq \sum_{i \in I} \boldsymbol{A}_{a,i} \leq \bar{\kappa}_a$. For each item i, we have lower and upper bounds on the total assignment of that item; $\underline{\psi}_i \leq \sum_{a \in N} \boldsymbol{A}_{a,i} \leq \bar{\psi}_i$. Finally, we have pairwise limits $\boldsymbol{C}_{a,i}$ for each agent a and item type i, requiring that $\boldsymbol{A}_{a,i} \leq \boldsymbol{c}_{a,i}$. It is always the case that the constraints define a finite set such that $|\mathcal{A}| \in \mathbb{N}$. In the example of reviewer assignment, these constraints reflect the review requirements per paper, load bounds for reviewers, conflicts of interest, and the constraint that no reviewer is assigned twice to any given paper.

Let $\mathbf{u} : \mathcal{A} \times \mathcal{V} \times \mathcal{G} \to \mathbb{R}$ be an affine function mapping from allocations to utilities for each group. $\mathbf{u}(\mathbf{a}, \mathbf{v}, G)$ denotes the utility of the group G under allocation $\mathbf{a} \in \mathcal{A}$ (recall that $\mathbf{a}$ is simply the vectorized version of the assignment $\boldsymbol{A}$). We assume that $\mathbf{u}$ is *additive*, such that the utility of group G under assignment $\mathbf{a} \in \mathcal{A}$ and valuations $\mathbf{v}$ is equal to $\frac{\mathbf{a}_G^\intercal \mathbf{v}_G}{|G|}$. For any group G, we will use the shorthand $\mathbf{u}_G$ to represent $\mathbf{u}(\mathbf{a}, \mathbf{v}, G)$. We then define a *welfare function* $W : \mathbb{R}^g \to \mathbb{R}$, where $W(\mathbf{u}_{G_1}, \mathbf{u}_{G_2}, \cdots, \mathbf{u}_{G_g})$ denotes the total welfare of allocation $\mathbf{a}$. The *weighted utilitarian social welfare* function is defined as $W^{\mathbf{w}}_{\text{USW}}(\mathbf{u}) \doteq \sum_{G=1}^{g} \boldsymbol{w}_G \mathbf{u}_G$, where $\boldsymbol{w} \in \mathbb{R}^g_{0+}$ denotes the weights on groups in $\mathcal{G}$. When $\boldsymbol{w}_G = |G|$ for all G, we call this function simply "utilitarian welfare" or "USW" and denote it as W_{USW}. The *group egalitarian social welfare* function (also "group egalitarian welfare" or "GESW") is defined as $W_{\text{GESW}}(\mathbf{u}) \doteq \min_{G=1}^{g} \mathbf{u}_G$. We do not consider individual egalitarian welfare in this work; under robust and stochastic optimization the egalitarian welfare is zero when the number of items is proportional to the number of agents and uncertainty is non-trivial.

2.2 Optimizing Allocations under Uncertainty

We consider two main approaches to dealing with uncertainty: the *robust optimization* approach and the *Conditional Value at Risk* approach.
In the robust approach, we obtain an *uncertainty set* $\mathcal{V}$ that contains the true agent valuations $\mathbf{v}^*$ with probability $\geq 1-\delta$ for some confidence parameter $\delta \in [0,1)$. We then optimize the welfare corresponding to the *worst* valuation matrix in the uncertainty set, i.e., $\max_{\mathbf{a}\in\mathcal{A}} \min_{\mathbf{v}\in\mathcal{V}} W(\mathbf{u}(\mathbf{a}))$. This approach is appropriate when we do not have access to a full distribution $\mathcal{D}_{\mathbf{v}}$ but have error bounds on $\mathbf{v}^*$ [14]. We investigate this regime in Section 3.
When we have access to a full distribution $\mathcal{D}_{\mathbf{v}}$ over a random variable $\mathbf{v} \in [0,1]^{nm}$, we apply a stochastic approach instead. We compute the welfare distribution and optimize the conditional expectation over an α-percentile of the welfare or *Conditional Value at Risk at* α (CVaR_α), where the confidence parameter α is determined by the market-maker. This approach is also referred to as the *soft-robust* approach. Suppose that $\mathcal{D}_{\mathbf{v}}$ represents the probability distribution of the random valuation matrix $\boldsymbol{V}$ and α denotes the percentile of the welfare we wish to optimize. For any $\alpha \in (0, 0.5)$, CVaR_α is defined as $\mathbb{E}_{\mathbf{v}\sim\mathcal{D}_{\mathbf{v}}}[X \mid X \leq \nu_\alpha(W;\mathbf{a},\mathbf{v})]$ where $\nu_\alpha(W;\mathbf{a},\mathbf{v})$ denotes the α-percentile of welfare. This approach is only appropriate when $\mathcal{D}_{\mathbf{v}}$ is fully known, and $\mathbb{E}_{V\sim\mathcal{D}_{\mathbf{v}}}[X \mid X \leq \nu_\alpha(W;\mathbf{a},\mathbf{v})]$ can be efficiently computed and optimized. We investigate this regime in Section 4.

Example 2.1 (The Importance of Considering Uncertainty)**.** Consider a simple two-agent, two-item instance, where each agent needs to get exactly one item, and either likes (utility 1) or dislikes it (utility 0). Agent preferences are Bernoulli random variables, where $\Pr[v_{1,1}=1]=0.8, \Pr[v_{1,2}=1]=0.9, \Pr[v_{2,1}=1]=0.5$, and $\Pr[v_{2,2}=1]=0.8$. If we maximize the sum of values over the expected value of each variable, we would assign i_1 to a_1 and i_2 to a_2, for a total expected value of 1.6. However, consider instead the objective of *Conditional Value at Risk*, which is the conditional expectation over the left tail of the distribution at a certain percentile. When we make the expectation-maximizing assignment, then $\Pr[W=0]=0.04$ and $\Pr[W=1]=0.32$. However, if we assign i_2 to agent a_1 and item i_1 to agent a_2, we have that $\Pr[W=0]=0.05$ and $\Pr[W=1]=0.5$. This means that the conditional expectation of welfare at the 30th percentile is higher if we assign i_2 to a_1 and i_1 to a_2 (it is .32 in the first case and .5 in the second case). If we want to retain welfare in the face of uncertainty, we might well choose to maximize this quantity rather than the expectation of the welfare.

3 Robust Welfare Optimization

We construct the optimization problems for Utilitarian and Egalitarian welfare objectives with the Robust approach. Many of these optimization problems are concave-convex max-min problems that can be directly solved using the naive projected gradient ascent technique [14]: in each iteration of the algorithm, the inner minimization problem is solved to optimality, followed by a projected-gradient step on the allocation $\mathbf{a}$. However, this method does not exploit the structure of these problems and is often computationally expensive or intractable, as demonstrated empirically in Section 5. Despite the inherent complexities of these problems, we show that under specific assumptions, these problems can be reduced to more manageable forms that are easier to optimize. We then discuss a range of algorithms for efficiently optimizing the simplified problems.

Scope: The Robust approach detailed in Section 2 assumes the availability of an uncertainty set of the valuation matrix. For the sake of computational tractability, we focus on the class of uncertainty sets defined by linear and ellipsoidal constraints:

$$\mathcal{V} = \left\{ \mathbf{v} \in \mathbb{R}^{nm} \mid \forall i \in [1,l],\, (\mathbf{v}-\bar{\mathbf{v}}_i)\Sigma_i^{-1}(\mathbf{v}-\bar{\mathbf{v}}_i) \leq r_i^2, \boldsymbol{Q}\mathbf{v} \succeq \mathbf{e}, \mathbf{v} \succeq 0 \right\} \quad ,$$

where the i^{th} ellipsoidal uncertainty set has center $\bar{\mathbf{v}}_i \in \mathbb{R}^{nm}_{0+}$, covariance matrix $\Sigma_i \in \mathbb{R}^{nm\times nm}$, with radius $r_i \in \mathbb{R}$, $\boldsymbol{Q} \in \mathbb{R}^{k\times nm}$, and $\mathbf{e} \in \mathbb{R}^k$. We will further assume that the covariance matrices corresponding to the ellipsoidal uncertainty sets are positive semi-definite. This limitation on the structure of uncertainty sets is not too restrictive; it is possible to construct such uncertainty sets for linear regression and logistic regression models using statistical bounds, as shown in appendix D. Moreover, in all of our methods where obtaining an integer allocation is either not feasible or computationally tractable, we relax the set of feasible integer assignments $\mathcal{A} \subseteq \mathbb{N}^{n\times m}$ to a set of feasible continuous allocations $\tilde{\mathcal{A}} \subseteq \mathbb{R}_{0+}^{n\times m}$. We obtain integer allocations by applying randomized rounding techniques to the continuous allocations. This conversion procedure is laid out in [14].

3.1 Robust Allocation for Utilitarian Welfare

We consider the problem of finding an allocation that optimizes the utilitarian welfare under the worst valuation matrix in the uncertainty set. We formulate the problem as:

$$\max_{\mathbf{a}\in\mathcal{A}} \min_{\mathbf{v}\in\mathcal{V}} \sum_{G\in\mathcal{G}} \boldsymbol{w}_G \cdot \mathbf{u}(\mathbf{a}, \mathbf{v}, G) \ , \tag{1}$$

where $\boldsymbol{w}_G \forall G \in \mathcal{G}$ represent the scalar weight corresponding to group G. The objective and constraints of the inner-minimization problem described in (1) are convex, which confirms that the inner-minimization problem is also convex. Furthermore, the problem is strictly feasible, which satisfies Slater's condition [11] for strong duality. Therefore, by taking the dual of the inner-minimization problem, we can simplify the problem in (1) into a single equivalent maximization problem. We provide the dual formation in Proposition 3.1.

In the dual, let $\boldsymbol{\beta} \in \mathbb{R}_{0+}^k$ be the dual variable corresponding to the linear constraints $\boldsymbol{Q}\mathbf{v} \succeq \mathbf{e}$, $\boldsymbol{\lambda} \in \mathbb{R}_{0+}^l$ be the dual variable associated with the ellipsoidal constraints, and $\boldsymbol{\xi} \in \mathbb{R}^{nm}$ be the variable that combines the primal variable $\mathbf{a}$ with the dual variable of the non-negativity constraint on $\mathbf{v}$ for variable elimination. Furthermore, we define a set of feasible $\boldsymbol{\xi}$ as $\Lambda = \mathcal{A} - \mathbb{R}_{0+}^{nm}$, which is Pareto-dominated by $\mathcal{A}$.

Proposition 3.1. *The problem in* (1) *is equivalent to solving*

$$\max_{\substack{\boldsymbol{\xi}\in\Lambda,\boldsymbol{\lambda}\in\mathbb{R}_{0+}^l,\\ \boldsymbol{\beta}\in\mathbb{R}_{0+}^k}} \mathbf{c}^\intercal \Sigma_{[L]}^{-1} \mathbf{d}^\intercal + \boldsymbol{\beta}^\intercal \mathbf{e} - \frac{1}{4}\|\mathbf{c}^\intercal \Sigma_{[L]}^{-1/2}\|_2^2 + \sum_{i=1}^{l} \boldsymbol{\lambda}_i \|\bar{\mathbf{v}}_i^\intercal \Sigma_i^{-1/2}\|_2^2 - \|\mathbf{d}^\intercal \Sigma_{[L]}^{-1/2}\|_2^2 - \sum_{i=1}^{l} \boldsymbol{\lambda}_i r_i^2 \ , \tag{2}$$

where $\mathbf{c} = (-\boldsymbol{\beta}^\intercal \boldsymbol{Q} + \boldsymbol{\xi})$ *and* $\mathbf{d} = \sum_{i=1}^{l} \boldsymbol{\lambda}_i \bar{\mathbf{v}}_i^\intercal \Sigma_i^{-1}$, *and* $\Sigma_{[L]} = \sum_{i=1}^{l} \boldsymbol{\lambda}_i \Sigma_i^{-1}$. *Let* $\boldsymbol{\xi}^*$ *be the optimal* $\boldsymbol{\xi}$ *in* (2). *Then, the optimal allocation* $\mathbf{a}^*$ *can be derived from* $\boldsymbol{\xi}^*$ *by solving the system of equations:*

$$G \in \mathcal{G} : \frac{\boldsymbol{w}_G}{|G|} \cdot \mathbf{a}_G \preceq \boldsymbol{\xi}_G^*, \, \mathbf{a} \in \mathcal{A} \quad ,$$

Proposition 3.1 shows that the optimal allocation for the problem in eq. (1) can be computed by first solving the concave cubic program in eq. (2) to obtain $\boldsymbol{\xi}^*$ and then deriving the optimal allocation $\mathbf{a}^*$ from $\boldsymbol{\xi}^*$ by solving a system of equations. Notably, the problem in eq. (2) is a single maximization problem with fewer variables and constraints as compared to the max-min problem in (1), making it simpler to solve. Additionally, when the valuation uncertainty set is polyhedral, the problem in (2) simplifies further into a linear program (LP) which can be solved efficiently using standard LP solvers like Gurobi [27]. We present this result in proposition 3.2. Moreover, when the valuation uncertainty set has a single ellipsoidal constraint with a non-negativity constraint, we can compute the exact optimal solution using iterated quadratic programming (IQP), as described in Proposition 3.3.

Proposition 3.2. *In the case where the uncertainty set* $\mathcal{V}$ *is defined purely by linear constraints, i.e.,* $\mathcal{V} = \{\mathbf{v} \in \mathbb{R}_{0+}^{nm} \mid \boldsymbol{Q}\mathbf{v} \succeq \mathbf{e}\}$, *the optimal allocation* $\mathbf{a}^*$ *for the problem in* (1) *can be computed by solving the linear program:*

$$\max_{\mathbf{a}\in\mathcal{A},\boldsymbol{\beta}\in\mathbb{R}_{0+}^k} \boldsymbol{\beta}^\intercal \mathbf{e} \quad \text{s.t. } \forall G \in \mathcal{G} : \ \boldsymbol{\beta}_G^\intercal \boldsymbol{Q}_G \preceq \frac{\boldsymbol{w}_G}{|G|} \mathbf{a}_G \ . \tag{3}$$

Proposition 3.3. *Suppose that the set* $\mathcal{V}$ *in* (1) *is defined by a single truncated ellipsoidal constraint i.e.,* $\mathcal{V} = \{\mathbf{v} \in \mathbb{R}_{0+}^{nm} \mid (\mathbf{v} - \bar{\mathbf{v}})\Sigma_i^{-1}(\mathbf{v} - \bar{\mathbf{v}}) \leq r^2\}$. *The problem in* (1) *is equivalent to solving*

$$\max_{\lambda\in\mathbb{R}_{0+},\boldsymbol{\xi}\in\Lambda} \left(\boldsymbol{\xi}^\intercal \bar{\mathbf{v}} - \frac{\|\boldsymbol{\xi}^\intercal \Sigma^{\frac{1}{2}}\|_F^2}{4\lambda} - \lambda r^2 \right) \ . \tag{4}$$

The exact optimal solution $(\lambda^*, \boldsymbol{\xi}^*)$ *to eq.* (4) *can be computed by alternately performing two steps until convergence: first, fixing* $\boldsymbol{\xi}$ *and optimizing* λ, *i.e.,* $\lambda = \|\boldsymbol{\xi}^\intercal \Sigma^{\frac{1}{2}}\|_F^2 / 2r$, *and second, fixing* λ *and solving a concave quadratic program to optimize* $\boldsymbol{\xi}$. *Furthermore, the optimal allocation* $\mathbf{a}^*$ *can be computed from* $\boldsymbol{\xi}^*$ *as in Proposition 3.1.*

3.2 Robust Allocation for Group Egalitarian Welfare

We now consider the problem where we aim to maximize the welfare corresponding to the worst group while using the robust approach for handling uncertainty. We can represent this problem as

$$\max_{\mathbf{a}\in\mathcal{A}} \min_{\mathbf{v}\in\mathcal{V}} \min_{G\in\mathcal{G}} \; \mathbf{u}(\mathbf{a},\mathbf{v},G) \;\; . \tag{5}$$

This problem presents inherent challenges due to the non-smoothness of the inner-minimization problem and the joint constraint on the uncertainties of the valuation matrices of different groups. These factors make it difficult to compute the dual and reduce the problem or efficiently solve the problem using the quadratic program technique described in Proposition 3.3. To streamline this problem, we assume that the uncertainty sets for each group $G \in \mathcal{G}$ are independent of each other. This allows us to represent the uncertainty set $\mathcal{V}$ as a Cartesian product of an individual group's uncertainty set, $\mathcal{V} = \bigotimes_{G\in\mathcal{G}} \mathcal{V}_G$. Furthermore, we can reorder the two inner-minimization problems without compromising generality.

$$\max_{\mathbf{a}\in\mathcal{A}} \min_{G\in\mathcal{G}} \min_{\mathbf{v}_G\in\mathcal{V}_G} \frac{1}{|G|}\mathbf{a}_G^\intercal \mathbf{v}_G \;\; . \tag{6}$$

We note that the problem in (6) is a concave-convex optimization problem that can be solved exactly using the sub-gradient ascent method.

An alternative approach to optimizing the problem (6) involves taking the dual of the inner-most minimization and reordering the inner-minimization over groups and the inner-maximization problem over the dual variables to obtain a single max-min problem. This simplified problem can then be solved with iterated max-min quadratic programming. We illustrate this result in Proposition 3.4.

Proposition 3.4. *The problem in* (5) *is equivalent to solving*

$$\max_{\substack{\boldsymbol{\xi}\in\Lambda,\\ \boldsymbol{\lambda}\in\mathbb{R}_{0+}^{g\times l},\\ \boldsymbol{\beta}\in\mathbb{R}_{0+}^{g\times k}}} \min_{G\in\mathcal{G}} \boldsymbol{\beta}_G^\intercal \mathbf{e}_G + \mathbf{c}_G^\intercal \Sigma_*^{-1}\mathbf{d}_G - \frac{1}{4}\|\mathbf{c}_G^\intercal \Sigma_*^{-1/2}\|_2^2 + \sum_{i=1}^{l}\left(\boldsymbol{\lambda}_{G,i}\|\bar{\mathbf{v}}_{G,i}^\intercal \Sigma_*^{-1/2}\|_2^2 - \boldsymbol{\lambda}_{G,i} r_{G,i}^2\right) - \|\mathbf{d}_G^\intercal \Sigma_*^{-1/2}\|_2^2 \tag{7}$$

and $\forall G \in \mathcal{G} : \; \mathbf{c}_G = (\boldsymbol{\xi}_G - \boldsymbol{\beta}_G^\intercal \boldsymbol{Q}_G)$, $\mathbf{d}_G = \sum_{i=1}^{l} \boldsymbol{\lambda}_{G,i}\bar{\mathbf{v}}_{G,i}\Sigma_{G,i}^{-1}$, *and* $\Sigma_* = \sum_{i=1}^{l}, \boldsymbol{\lambda}_{G,i}\Sigma_{G,i}^{-1}$. *Let* $\boldsymbol{\xi}^*$ *be the optimal* $\boldsymbol{\xi}$ *in* (7)*. Then, the optimal allocation* $\mathbf{a}^*$ *satisfies the system of equations:*

$$G \in \mathcal{G} : \frac{\boldsymbol{w}_G}{|G|} \cdot \mathbf{a}_G \preceq \boldsymbol{\xi}_G^*, \; \mathbf{a} \in \mathcal{A} \;\; ,$$

The dual variables $\boldsymbol{\lambda}_G, \boldsymbol{\beta}_G, \boldsymbol{\zeta}_G$ and $\boldsymbol{\xi}_G$ for each group G are interpreted as in Proposition 3.1. The optimization problem in (7) is concave with respect to the dual variables $\boldsymbol{\lambda}, \boldsymbol{\beta}$ and $\boldsymbol{\xi}$. Consequently, we can solve it using an approach similar to that in proposition 3.3. Specifically, we employ the max-min iterated quadratic programming [44], alternately fixing $\boldsymbol{\lambda}$ and optimizing the rest of the dual variables $(\boldsymbol{\beta}, \boldsymbol{\xi})$ and vice versa until convergence.

Interestingly, even when optimizing the egalitarian welfare objective with only polyhedral uncertainty sets, the robust egalitarian problem described in (7) simplifies to a straightforward linear program. This is akin to what we observe in the robust utilitarian case (Proposition 3.2). We formalize this finding in Proposition 3.5.

Proposition 3.5. *In the case where the uncertainty set* $\mathcal{V}$ *is defined only by linear constraints, i.e.,* $\mathcal{V} = \{\mathbf{v} \in \mathbb{R}^{nm} \mid \boldsymbol{Q}\mathbf{v} \succeq \mathbf{e}, \mathbf{v} \succeq 0\}$*, the max-min-min problem in* (5) *is trivially transformable into a linear program.*

3.3 Robust Allocation for Monotonic Welfare Functions

We now extend our findings to a broader class of monotonic welfare functions. Specifically, we show that when optimizing a monotonic welfare objective under the assumption that valuation uncertainty sets across groups are independent, we can decompose the problem into sub-problems such that we independently determine the worst valuation in the uncertainty set of each group, while jointly optimizing the allocations of different groups.

Proposition 3.6. *Consider an optimization problem of the form*

$$\max_{\mathbf{a}\in\mathcal{A}} \min_{\mathbf{v}\in\mathcal{V}} W_M(\mathbf{u}(\mathbf{a},\mathbf{v},G_1), \mathbf{u}(\mathbf{a},\mathbf{v},G_2), \ldots, \mathbf{u}(\mathbf{a},\mathbf{v},G_g)) \;\; , \tag{8}$$

where the welfare function $W_M(\cdot)$ is monotonic in the utility of groups. If the valuation uncertainty sets are independent across groups, $\mathcal{V} = \bigotimes_{G\in\mathcal{G}} \mathcal{V}_G$, then, the problem in (8) *simplifies to*

$$\max_{\mathbf{a}\in\mathcal{A}} W_M(\min_{\mathbf{v}\in\mathcal{V}} \mathbf{u}(\mathbf{a},\mathbf{v},G_1), \min_{\mathbf{v}\in\mathcal{V}} \mathbf{u}(\mathbf{a},\mathbf{v},G_2), \ldots, \min_{\mathbf{v}\in\mathcal{V}} \mathbf{u}(\mathbf{a},\mathbf{v},G_g)) \ .$$

We note that the egalitarian problem in (5) is an instance of the class of optimization problem described in (8). Furthermore, when the allocation and valuation uncertainty sets are convex and compact, the problem in (8) can be solved using constrained convex-concave minimax optimization algorithms [16, 24, 55].

4 Stochastic Welfare Optimization

When we have a distribution $\mathcal{D}_{\mathbf{v}}$ over the valuation matrix, we leverage the CVaR measure to compute high-confidence robust allocations for utilitarian and egalitarian welfare objectives.

4.1 CVaR Allocation for Utilitarian Welfare

We wish to find an allocation that maximizes the CVaR_α of the utilitarian welfare. Let $\tilde{\mathbf{v}}$ represent the random valuation vector. Then, for any confidence level α, we can formulate this problem as:

$$\max_{\mathbf{a}\in\mathcal{A}} \mathrm{CVaR}_\alpha\left[\sum_{G\in\mathcal{G}} w_G\cdot\mathbf{u}(\mathbf{a},\tilde{\mathbf{v}},G)\right] \equiv \max_{\mathbf{a}\in\mathcal{A},b\in\mathbb{R}}\left\{b-\frac{1}{\alpha}\mathop{\mathbb{E}}_{\mathbf{v}\sim\mathcal{D}_{\mathbf{v}}}\left[b-\left(\sum_{G\in\mathcal{G}}\frac{\boldsymbol{w}_G}{|G|}\cdot\mathbf{a}_G^\intercal\mathbf{v}_G\right)_+\right]\right\} \ , \tag{9}$$

where $(x)_+ = \max(x,0)$ represents the positive part of x. Computing the exact expectation in this problem may not be feasible for every distribution $\mathcal{D}_{\mathbf{v}}$. Therefore, we adopt a sampling-based approach to approximately optimize the CVaR of utilitarian welfare. We begin by drawing h samples of the valuation matrix from $\mathcal{D}_{\mathbf{v}}$ represented as $\mathbf{v}^1,\mathbf{v}^2,\mathbf{v}^3,\ldots\mathbf{v}^h$. We then use these samples to solve the problem described in (9) by solving the linear program outlined in Proposition 4.1.

Proposition 4.1. *Given h samples of $\tilde{\mathbf{v}}$, i.e., $\mathbf{v}^1,\mathbf{v}^2,\mathbf{v}^3,\ldots\mathbf{v}^h$ from $\mathcal{D}_{\mathbf{v}}$, the optimal allocation for the problem in* (9) *can be approximately computed by solving*

$$\max_{\mathbf{a}\in\mathcal{A}}\ \max_{\boldsymbol{y}\in\mathbb{R}^h_{0+},b\in\mathbb{R}}\left(b-\frac{1}{\alpha}\sum_{j=1}^h\boldsymbol{y}_j\right)\qquad\forall j\in[1,h]:\ \boldsymbol{y}_j\geq\frac{1}{h}\left(b-\sum_{G\in\mathcal{G}}\frac{\boldsymbol{w}_G}{|G|}\cdot\mathbf{a}_G^\intercal\mathbf{v}_G^j\right) \ . \tag{10}$$

The CVaR estimator used in (10) is a strongly consistent estimator [28]. Therefore, the approximation error of the objective in (10) goes to 0 as $h\to\infty$. In proposition 4.3, we bound the sample complexity of the problem in (10) when the valuation matrix is sub-Gaussian distributed.

Assumption 4.2. *Let $\tilde{w}$ represent the random welfare corresponding to a given allocation $\mathbf{a}$ and let f be its density function. Furthermore, let $\nu_\alpha(W;\mathbf{a},\mathbf{v})$ denotes the α-percentile of the welfare corresponding to allocation $\mathbf{a}$. There exists universal constants $\eta,\delta'\geq 0$, s.t., $f(w)\geq\eta\,\forall w\in[\nu_\alpha(W;\mathbf{a},\mathbf{v})-\frac{\delta'}{2},\nu_\alpha(W;\mathbf{a},\mathbf{v})+\frac{\delta'}{2}]$.*

For any allocation $\mathbf{a}$, let $\hat{c}_{h,\alpha}(\mathbf{a})$ represent the empirical estimate of CVaR of utilitarian welfare computed from h samples and $c_{h,\alpha}(\mathbf{a})$ represent the corresponding true value.

Proposition 4.3. *Suppose that $\mathbf{v}$ is a multivariate sub-Gaussian with mean $\bar{\mathbf{v}}\in\mathbb{R}^{nm}$ and covariance proxy $\Sigma\in\mathbb{R}^{nm\times nm}$, i.e., $\exists K\geq 0$ s.t. $\mathbb{E}\left[\exp(\lambda(\mathbf{v}-\bar{\mathbf{v}})^\intercal\boldsymbol{z}))\right]\leq\exp(\lambda^2K^2\boldsymbol{z}^\intercal\Sigma\boldsymbol{z}/2),\forall\lambda\in\mathbb{R},\forall\boldsymbol{z}\in\mathbb{R}^{nm}$ and that Assumption 4.2 holds. Let $|\mathcal{A}|$ represent the number of feasible allocations, $\mathbf{a}'_G=\frac{\boldsymbol{w}_G}{|G|}\cdot\mathbf{a}_G$, and $h=\mathcal{O}\left(\frac{8\max(\max_{\mathbf{a}\in\mathcal{A}}\mathbf{a}'^\intercal\Sigma\mathbf{a}',8)\log\left(\frac{6|\mathcal{A}|}{\delta}\right)}{\varepsilon^2(\alpha)^2\min(\eta^2,1)}\right)$ where $\delta\in(0,1)$. Then, $\Pr[\forall\mathbf{a}\in\mathcal{A}:|\hat{c}_{h,\alpha}(\mathbf{a})-c_\alpha(\mathbf{a})|\leq\varepsilon]\geq 1-\delta$.*

When the valuation $\mathbf{v}$ is normally distributed, we can circumvent the sampling approach and instead solve the problem directly by optimizing a quadratic optimization problem (Proposition 4.4), which depends solely on the mean and covariance of the valuation matrix $\mathbf{v}$.

Proposition 4.4. *If the valuation $\mathbf{v}$ is distributed as a multivariate Gaussian, i.e., $\mathbf{v}\sim\mathcal{N}(\bar{\mathbf{v}},\Sigma)$, then, the optimization problem in* (9) *simplifies to*

$$\max_{\mathbf{a}\in\mathcal{A}}\mathbf{a}^\intercal\bar{\mathbf{v}}-\frac{\phi(\Phi^{-1}(1-\alpha))}{\alpha}\sqrt{\mathbf{a}^\intercal\Sigma\mathbf{a}} \ . \tag{11}$$

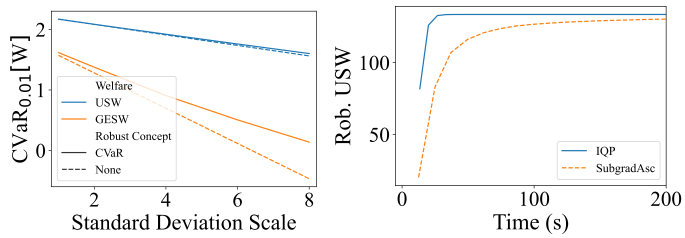

Figure 1: Left: CVaR as noise increases for AAMAS 2015. Right: Convergence behavior of the Iterated Quadratic Program (IQP) vs. Subgradient Ascent (SA) approach on AAMAS 2015.

The problem in (11) is concave and can be solved exactly using the projected gradient ascent method.

4.2 CVaR Allocation for Group Egalitarian Welfare

For our final objective, we wish to optimize egalitarian welfare under uncertainty using the CVaR approach. We formulate this optimization problem as

$$\max_{\mathbf{a}\in\mathcal{A}} \mathrm{CVaR}_\alpha \left[\min_{G\in\mathcal{G}} \mathbf{u}(\mathbf{a},\tilde{\mathbf{v}},G)\right] \equiv \max_{\mathbf{a}\in\mathcal{A}, w\in\mathbb{R}} \left\{ w - \frac{1}{\alpha}\mathbb{E}\left[\left(w - \min_{G\in\mathcal{G}} \frac{1}{|G|}\cdot \mathbf{a}_G^\intercal \tilde{\mathbf{v}}_G\right)_+\right]\right\} . \tag{12}$$

To optimize the problem described in (12), we solve a linear program similar to the one used for optimizing the CVaR Utilitarian objective in (9). We refer the readers to Proposition H.3 for more details. When the valuation matrix $\mathbf{v}$ is normally distributed and the uncertainty sets of different groups are independent, the result is a quadratic program characterized by a linear objective and quadratic constraints, as detailed in proposition C.1.

5 Experiments

We run experiments on three reviewer assignment datasets. The datasets contain bids from the International Conference on Autonomous Agents and Multiagent Systems (AAMAS) 2015, 2016, and 2021 [40, 41].[1] We consider the papers as the "agents" and the reviewers as the "items." This is a fairly standard assumption in most recent reviewer assignment approaches, reflecting the primary goal of peer review to assign qualified and interested reviewers to papers [14, 29, 30, 37, 45, 54].

Reviewers issue bids of `yes`, `maybe`, `no`, or `no response`. We run two experiments with this data. In one, we binarize the bids such that `yes` and `maybe` are considered affirmative and `no` is considered negative, while in the other we convert the bids to numerical scores such that `yes` is 1, `maybe` is .5, and `no` is 0.01. Under the binarized model, we fit a logistic matrix factorization model to predict whether the bid is affirmative or negative, and in the continuous model, we fit a Gaussian process matrix factorization model [35]. We derive probability distributions and uncertainty sets from these models. More details on prediction and uncertainty set construction are in Appendix E. These datasets do not contain groups of papers and reviewers, so we create 4 roughly balanced clusters of reviewers and papers for each dataset using the procedure outlined in Appendix F. We define our valid set of assignments $\mathcal{A}$ as follows. For each paper $a \in N$, we set $\underline{\kappa}_a = \bar{\kappa}_a = 3$ for all a in AAMAS 2015, and $\underline{\kappa}_a = \bar{\kappa}_a = 2$ for all a in AAMAS 2016 and 2021. For each reviewer i, we set $\underline{\psi}_i = 0$ and $\bar{\psi}_i = 15$ for 2015 and 2016 and 4 for 2021. We optimize and evaluate $\mathrm{CVaR}_{0.01}$; we take $4,000$ samples from the distribution to optimize for CVaR using the sampling-based approach, and we take $10,000$ samples to estimate the CVaR for evaluation. We optimize and evaluate the adversarial welfares at the $\delta = 0.3$ level (there is a 70% chance the true values lie in the uncertainty set). All results are averaged over 5 runs of subsampling 20% of each dataset.

[1] Available at https://preflib.simonrey.fr/dataset/00037.

Table 2: Performance of different allocations across each metric on the AAMAS 2015 dataset.

Allocation	Evaluation Objective					
	USW	GESW	CVaR USW	CVaR GESW	Rob. USW	Rob. GESW
USW	$\mathbf{1.00 \pm 0}$	$\mathbf{1.00 \pm 0}$	$\mathbf{1.00 \pm 0}$	$\mathbf{1.00 \pm 0}$	0 ± 0	0 ± 0
GESW	0.97 ± 0.01	$\mathbf{1.00 \pm 0}$	0.97 ± 0.01	0.97 ± 0.02	0 ± 0	0 ± 0
CVaR USW	$\mathbf{1.00 \pm 0}$	0.99 ± 0	$\mathbf{1.00 \pm 0}$	0.99 ± 0	0 ± 0	0 ± 0
CVaR GESW	0.98 ± 0	0.99 ± 0	0.97 ± 0.01	$\mathbf{1.00 \pm 0}$	0 ± 0	0 ± 0
Rob. USW	0.92 ± 0.01	0.90 ± 0.02	0.92 ± 0.01	0.90 ± 0.02	$\mathbf{1.00 \pm 0}$	$\mathbf{1.00 \pm 0}$
Rob. GESW	0.89 ± 0.04	0.85 ± 0.06	0.89 ± 0.04	0.86 ± 0.06	0.88 ± 0.02	$\mathbf{1.00 \pm 0}$

5.1 Results

Overall Performance Section 5.1 shows the results for the binarized version of AAMAS 2015 bids. Similar tables for the 5 other settings are included in Appendix G. Each row shows the metrics for the allocation produced by the method which optimizes for the objective shown in the left-most column. Objective values are normalized by dividing by the maximum value of that objective per seed. All methods have 0 adversarial welfare, even at the $\delta = 0.3$ level, indicating that if robustness to adversarial noise is desired, it is very important to consider this objective explicitly. We approximate the optimal $\mathrm{CVaR}_{0.01}$ using 1000 samples, which leaves some room for sampling error as evidenced by the strong performance of the baseline USW and GESW allocations on the $\mathrm{CVaR}_{0.01}$ measure. However, relatively little noise is actually present in this dataset, as the $\mathrm{CVaR}_{0.01}$ is relatively high for both USW and GESW in all cases.

Robustness under Increasing Uncertainty Figure 1 shows the $\mathrm{CVaR}_{0.01}$ on the Gaussian version of all three datasets as we artificially increase the amount of noise. We multiply the standard deviations of the Gaussian distributions by a scalar and optimize for the CVaR or the central estimate of the USW and GESW. We then plot $\mathrm{CVaR}_{0.01}$ as noise increases. Although the CVaR approach is less important at low noise levels, the CVaR of welfare decreases for both welfare measures as noise increases. GESW has a sharper decline than USW. We see that as the noise increases, the $\mathrm{CVaR}_{0.01}$ of the baseline USW and GESW maximizing allocations drops off relative to the same value for the CVaR-optimized allocation.

Runtime Finally, for the soft robust optimization setting with ellipsoidal uncertainty sets (derived from confidence intervals over the Gaussian process matrix factorization), we compare the IQP approach Proposition 3.3 to projected sub-gradient ascent on the original max-min problem. We find that IQP converges much faster than the subgradient ascent algorithm; see Figure 3. Sub-gradient ascent fails to converge in $1,000$ iterations for the adversarial GESW objective on all datasets and the USW objective on AAMAS 2021.

6 Limitations and Conclusion

The CVaR approach requires solving linear programs with a large number of samples to be effective, which makes them computationally expensive. One potential solution is to leverage importance sampling methods to reduce the variance of the estimator [17, 58]. Furthermore, future research could benefit from empirically and theoretically analyzing other fairness objectives like Nash Welfare [13], Gini Index [23], and Envy-Freeness [38].

In conclusion, we explore the stochastic and robust optimization regimes for utilitarian and group-wise welfare objectives. The robust optimization algorithms depend on the form of the uncertainty set. We show that when the uncertainty set has linear constraints only, the resulting problem is an LP and can be solved efficiently. Under ellipsoidal constraints, we demonstrate an iterative quadratic programming approach converges much faster than the naive subgradient approach for utilitarian welfare. However, the robust egalitarian welfare remains challenging to optimize. In the stochastic regime, we lay out the sample complexity of CVaR for the utilitarian welfare objective. We demonstrate the feasibility of estimating probability distributions and uncertainty sets on three years of bid data from AAMAS, and show that the robust and CVaR approaches laid out in this paper combat the uncertainty present in these three datasets.

References

[1] Emily L Aiken, Guadalupe Bedoya, Joshua E Blumenstock, and Aidan Coville. Program targeting with machine learning and mobile phone data: Evidence from an anti-poverty intervention in afghanistan. *Journal of Development Economics*, 161:103016, 2023.

[2] Amine Allouah, Christian Kroer, Xuan Zhang, Vashist Avadhanula, Anil Dania, Caner Gocmen, Sergey Pupyrev, Parikshit Shah, and Nicolas Stier. Robust and fair work allocation. *arXiv preprint arXiv:2202.05194*, 2022.

[3] Igor Averbakh. Minmax regret solutions for minimax optimization problems with uncertainty. *Operations Research Letters*, 27(2):57–65, 2000.

[4] Haris Aziz, Evi Micha, and Nisarg Shah. Group fairness in peer review. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2023. Extended Abstract, to Appear.

[5] Mohammad Javad Azizi, Phebe Vayanos, Bryan Wilder, Eric Rice, and Milind Tambe. Designing fair, efficient, and interpretable policies for prioritizing homeless youth for housing resources. In *Integration of Constraint Programming, Artificial Intelligence, and Operations Research: 15th International Conference, CPAIOR 2018, Delft, The Netherlands, June 26–29, 2018, Proceedings 15*, pages 35–51. Springer, 2018.

[6] Aharon Ben-Tal, Laurent El Ghaoui, and Arkadi Nemirovski. *Robust optimization*, volume 28. Princeton university press, 2009.

[7] Dimitris Bertsimas, David B Brown, and Constantine Caramanis. Theory and applications of robust optimization. *SIAM review*, 53(3):464–501, 2011.

[8] John R Birge and Francois Louveaux. *Introduction to stochastic programming*. Springer Science & Business Media, 2011.

[9] Craig Boutilier, Relu Patrascu, Pascal Poupart, and Dale Schuurmans. Constraint-based optimization and utility elicitation using the minimax decision criterion. *Artificial Intelligence*, 170(8-9):686–713, 2006.

[10] Craig Boutilier. Computational decision support: Regret-based models for optimization and preference elicitation, 2013.

[11] Stephen Boyd, Stephen P Boyd, and Lieven Vandenberghe. *Convex optimization*. Cambridge University Press, 2004.

[12] Jan Buermann, Enrico H Gerding, and Baharak Rastegari. Fair allocation of resources with uncertain availability. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*, pages 204–212, 2020.

[13] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D Procaccia, Nisarg Shah, and Junxing Wang. The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation (TEAC)*, 7(3):1–32, 2019.

[14] Cyrus Cousins, Justin Payan, and Yair Zick. Into the unknown: Assigning reviewers to papers with uncertain affinities. *Proceedings of the XXth International Symposium on Algorithmic Game Theory (SAGT)*, 2023.

[15] Cyrus Cousins. An axiomatic theory of provably-fair welfare-centric machine learning. In *Proceedings of the 35th Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 16610–16621, 2021.

[16] Constantinos Daskalakis, Stratis Skoulakis, and Manolis Zampetakis. The complexity of constrained min-max optimization. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2021, page 1466–1478, New York, NY, USA, 2021. Association for Computing Machinery.

[17] Anand Deo and Karthyek Murthy. Efficient black-box importance sampling for var and cvar estimation. In *2021 Winter Simulation Conference (WSC)*, pages 1–12, 2021.

[18] Siddartha Devic, David Kempe, Vatsal Sharan, and Aleksandra Korolova. Fairness in matching under uncertainty. *Proceedings of the 40th International Conference on Machine Learning (ICML)*, 2023.

[19] Steven Diamond and Stephen Boyd. Cvxpy: a python-embedded modeling language for convex optimization. *J. Mach. Learn. Res.*, 17(1):2909–2913, jan 2016.

[20] Kate Donahue and Jon Kleinberg. Fairness and utilization in allocating resources with uncertain demand. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*, pages 658–668, 2020.

[21] Evan Dong and Cyrus Cousins. Decentering imputation: Fair learning at the margins of demographics. In *Queer in AI Workshop @ ICML*, 2022.

[22] John C Duchi and Hongseok Namkoong. Learning models with uniform performance via distributionally robust optimization. *The Annals of Statistics*, 49(3):1378–1406, 2021.

[23] Corrado Gini. On the measure of concentration with special reference to income and statistics. *Colorado College Publication, General Series*, 208(1):73–79, 1936.

[24] Denizalp Goktas and Amy Greenwald. Convex-concave min-max stackelberg games. *Advances in Neural Information Processing Systems*, 34:2991–3003, 2021.

[25] Bram L Gorissen, İhsan Yanıkoğlu, and Dick Den Hertog. A practical guide to robust optimization. *Omega*, 53:124–137, 2015.

[26] Negin Gorlezaei, Patrick Jaillet, and Zijie Zhou. Online resource allocation with samples. *arXiv preprint arXiv:2210.04774*, 2022.

[27] Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, 2023.

[28] L. Jeff Hong and Guangwu Liu. Monte carlo estimation of value-at-risk, conditional value-at-risk and their sensitivities. In *Proceedings of the Winter Simulation Conference*, WSC '11, page 95–107. Winter Simulation Conference, 2011.

[29] Steven Jecmen, Hanrui Zhang, Ryan Liu, Nihar Shah, Vincent Conitzer, and Fei Fang. Mitigating manipulation in peer review via randomized reviewer assignments. In *Proceedings of the 34th Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 12533–12545, 2020.

[30] Ari Kobren, Barna Saha, and Andrew McCallum. Paper matching with local fairness constraints. In *Proceedings of the 25th International Conference on Knowledge Discovery and Data Mining (KDD)*, pages 1247–1257, 2019.

[31] Panos Kouvelis and Gang Yu. *Robust discrete optimization and its applications*, volume 14. Springer Science & Business Media, 2013.

[32] Pavlo Krokhmal, Jonas Palmquist, and Stanislav Uryasev. Portfolio optimization with conditional value-at-risk objective and constraints. *Journal of risk*, 4:43–68, 2002.

[33] Amanda R Kube, Sanmay Das, and Patrick J Fowler. Fair and efficient allocation of scarce resources based on predicted outcomes: implications for homeless service delivery. *Journal of Artificial Intelligence Research*, 76:1219–1245, 2023.

[34] Prashanth L.A., Krishna Jagannathan, and Ravi Kolla. Concentration bounds for CVaR estimation: The cases of light-tailed and heavy-tailed distributions. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 5577–5586. PMLR, 13–18 Jul 2020.

[35] Neil D Lawrence and Raquel Urtasun. Non-linear matrix factorization with gaussian processes. In *Proceedings of the 26th annual international conference on machine learning*, pages 601–608, 2009.

[36] Daniel Levy, Yair Carmon, John C Duchi, and Aaron Sidford. Large-scale methods for distributionally robust optimization. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 8847–8860. Curran Associates, Inc., 2020.

[37] Kevin Leyton-Brown, Yatin Nandwani, Hedayat Zarkoob, Chris Cameron, Neil Newman, Dinesh Raghu, et al. Matching papers and reviewers at large conferences. *Artificial Intelligence*, 331:104119, 2024.

[38] Richard J Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC)*, pages 125–131, 2004.

[39] Elita A Lobo, Mohammad Ghavamzadeh, and Marek Petrik. Soft-robust algorithms for batch reinforcement learning. *arXiv preprint arXiv:2011.14495*, 2020.

[40] Nicholas Mattei and Toby Walsh. Preflib: A library of preference data HTTP://PREFLIB.ORG. In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT 2013)*, Lecture Notes in Artificial Intelligence. Springer, 2013.

[41] Nicholas Mattei and Toby Walsh. A preflib.org retrospective: Lessons learned and new directions. *Trends in Computational Social Choice. AI Access Foundation*, pages 289–309, 2017.

[42] Hongseok Namkoong and John C Duchi. Stochastic gradient methods for distributionally robust optimization with f-divergences. *Advances in neural information processing systems*, 29, 2016.

[43] Włodzimierz Ogryczak, Adam Wierzbicki, and Marcin Milewski. A multi-criteria approach to fair and efficient bandwidth allocation. *Omega*, 36(3):451–463, 2008. Special Issue on Multiple Criteria Decision Making for Engineering.

[44] A.R. Pankov, E.N. Platonov, and Konstantin Siemenikhin. Minimax quadratic optimization and its application to investment planning. *Automation and Remote Control*, 62:1978–1995, 01 2001.

[45] Justin Payan and Yair Zick. I will have order! Optimizing orders for fair reviewer assignment. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 440–446, 2022.

[46] Dominik Peters, Ariel D Procaccia, and David Zhu. Robust rent division. In *Proceedings of the 36th Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2022.

[47] András Prékopa. *Stochastic programming*, volume 324. Springer Science & Business Media, 2013.

[48] David Pujol, Ryan McKenna, Satya Kuppam, Michael Hay, Ashwin Machanavajjhala, and Gerome Miklau. Fair decision making using privacy-protected data. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pages 189–199, 2020.

[49] Aida Rahmattalabi, Phebe Vayanos, Kathryn Dullerud, and Eric Rice. Learning resource allocation policies from observational data with an application to homeless services delivery. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 1240–1256, 2022.

[50] R Tyrrell Rockafellar, Stanislav Uryasev, et al. Optimization of conditional value-at-risk. *Journal of risk*, 2:21–42, 2000.

[51] Erik Rolland, Raymond A Patterson, Keith Ward, and Bajis Dodin. Decision support for disaster management. *Operations Management Research*, 3:68–79, 2010.

[52] Andrzej Ruszczyński and Alexander Shapiro. Stochastic programming models. *Handbooks in operations research and management science*, 10:1–64, 2003.

[53] Isaac Slavitt. Prioritizing municipal lead mitigation projects as a relaxed knapsack optimization: a method and case study. *International Transactions in Operational Research*, 30(6):3719–3737, 2023.

[54] Ivan Stelmakh, Nihar B Shah, and Aarti Singh. PeerReview4All: Fair and accurate reviewer assignment in peer review. In *Proceedings of the 30th International Conference on Algorithmic Learning Theory (ALT)*, pages 828–856, 2019.

[55] Yuanhao Wang and Jian Li. Improved algorithms for convex-concave minimax optimization. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 4800–4810. Curran Associates, Inc., 2020.

[56] Felix Wex, Guido Schryen, and Dirk Neumann. A fuzzy decision support model for natural disaster response under informational uncertainty. *International Journal of Information Systems for Crisis Response and Management (IJISCRAM)*, 4(3):23–41, 2012.

[57] Felix Wex, Guido Schryen, Stefan Feuerriegel, and Dirk Neumann. Emergency response in natural disaster management: Allocation and scheduling of rescue units. *European Journal of Operational Research*, 235(3):697–708, 2014.

[58] Tony Sit Yue Xing and Hoi Ying Wong. Variance reduction for risk measures with importance sampling in nested simulation. *Quantitative Finance*, 22(4):657–673, 2022.

[59] Wenbin Zhang and Jeremy C Weiss. Fair decision-making under uncertainty. In *2021 IEEE international conference on data mining (ICDM)*, pages 886–895. IEEE, 2021.

A Additional Related Work

Gorissen *et al.* [25] provide an excellent overview of optimization under uncertainty, including techniques used in this work, while Ben-Tal *et al.* [6], Bertsimas *et al.* [7] offer additional background on robust optimization. A standard approach in this regime is analyzing the dual of the uncertainty, as we generally do in this work. Stochastic optimization has a wide literature; the books by Birge and Louveaux [8], Levy *et al.* [36], Prékopa [47], Ruszczyński and Shapiro [52] present wide-ranging introductions to the topic. One of the primary approaches to stochastic optimization is conditional value at risk (CVaR), which can often be approximately optimized by sampling and optimizing over an objective composing the different samples [34, 39, 50]. We take this approach in our paper.

B Broader Impacts

We believe this work has the potential for a significant positive societal impact. Fair resource allocation algorithms are essential for various systems, including assigning reviewers in peer review processes, allocating resources to homeless and low-income populations, distributing emergency response resources during natural disasters, and resettling refugees. In this work, we develop methods for efficiently optimizing allocations of constrained resources under various fairness objectives while addressing uncertainty in resource preferences. These methods can be directly applied to the aforementioned problems. However, we advise users to conduct extensive testing on similar datasets before deploying these algorithms in real-world scenarios.

C Gaussian CVaR for Egalitarian Welfare

Proposition C.1. *If $\mathbf{v}_1, \mathbf{v}_2, \cdots, \mathbf{v}_g$ are $i.i.d$ and normally distributed, i.e., $\forall G \in \mathcal{G}, \mathbf{v}_G \sim \mathcal{N}(\bar{\mathbf{v}}_G, \Sigma_G)$, then, the optimization problem in* (12) *simplifies to*

$$\begin{aligned}
&\max_{\mathbf{a}\in\mathcal{A}, t\in\mathbb{R}} t \\
&\text{s.t. } \forall G \in \mathcal{G}: \left(\frac{1}{|G|}\cdot \mathbf{a}_G^\intercal \bar{\mathbf{v}}_G - t\right)^2 \geq \left(\frac{1}{|G|}\cdot\frac{\phi(\Phi^{-1}(1-\alpha))}{(\alpha)}\right)^2 \mathbf{a}_G^\intercal \Sigma_G \mathbf{a}_G \\
&\forall G \in \mathcal{G}: \left(\frac{1}{|G|}\cdot \mathbf{a}_G^\intercal \bar{\mathbf{v}}_G - t\right) \geq 0 \ .
\end{aligned} \tag{13}$$

The problem in (13) is a second order conic program (SOCP) and can be solved using popular SOCP solvers in CVXPY library [19].

D Constructing uncertainty sets

In this section we demonstrate a simple and natural approach to construct an uncertainty set using a logistic regression estimator. Logistic regression models with bounded cross-entropy loss result in polyhedral uncertainty sets. Replacing the logistic regression model with a model with bounded squared-error loss, or simply taking the confidence interval of a multivariate Gaussian, results in truncated ellipsoidal uncertainty sets. We construct uncertainty sets per group in all cases.

Assume we have a discrete set of c values $L \subseteq \mathbb{R}$, with $L = \{l_1, \dots l_c\}$. For each agent i and item type j we denote the true distribution over values $p^*(l|(i,j))$ and the distribution predicted by the logistic regression model is $\hat{p}(l|(i,j))$.

We estimate the cross-entropy loss of the model on a test set T, where $|T| = t$. This test set can be segmented by the group identity of the agent, such that we have $T_{G_1}, T_{G_2}, \dots T_{G_g}$ for each of the g groups (with sizes $t_{G_1}, \dots t_{G_g}$). We assume that the test set comes from the same distribution as the agent-item pairs of the assignment problem; this can be achieved either during dataset construction or by limiting the assignments (through the C constraints) to better reflect the test distribution. We can also apply likelihood reweighting in our uncertainty set construction, as in [14], though we do not do so here.

For an agent a and item type i, the cross-entropy loss of the distribution $\hat{p}$ with respect to the distribution p is defined as

$\mathbb{H}(p(l|(a,i)), \hat{p}(l|(a,i))) \doteq -\sum_{l \in L} p(l|(a,i)) \log \hat{p}(l|(a,i))$. For each T_G, we compute the mean of the cross-entropy loss $\hat{\xi}_G = \frac{1}{t_G} \sum_{(a,i) \in T_G} \mathbb{H}(p(l|(a,i)), \hat{p}(l|(a,i)))$, as well as the standard error of the mean $\hat{\eta}_G = \left(\frac{1}{t_G} \sum_{(i,j) \in T_G} (\mathbb{H}(p(l|(a,i)), \hat{p}(l|(a,i))) - \hat{\xi}_G)^2 \right)^{\frac{1}{2}}$. We model the distribution over cross-entropy losses for group G as $\mathcal{N}(\hat{\xi}_G, \hat{\eta}_G)$. We want an uncertainty set $\mathcal{V}$ such that the true values lie outside $\mathcal{V}$ with probability at most δ. Thus, using a union bound, we require each uncertainty set $\mathcal{V}_G$ for individual groups to contain the true valuations with probability at least $1 - \frac{\delta}{g}$. We can thus give the bound that the cross entropy loss is at most $\Phi^{-1}(1 - \frac{\delta}{g}, \hat{\xi}_G, \hat{\eta}_G)$, where $\Phi^{-1}(p, \mu, \sigma)$ denotes the p percentile of a normal distribution with mean μ and standard deviation σ.

In our assignment problem, for each group G with agents N_G we obtain the uncertainty set

$$\frac{1}{t_G m} \sum_{a \in N_G, i \in I} \mathbb{H}(p(l|(a,i)), \hat{p}(l|(a,i))) \leq \Phi^{-1}(1 - \frac{\delta}{g}, \hat{\xi}_G, \hat{\eta}_G) \ .$$

The bound can be made tighter if we restrict some pairs using C, in which case the cross-entropy term on the left side is only averaged over the pairs which are not restricted.

E Logistic and Gaussian Process Matrix Factorization

Both models define probability distributions over outcomes, which we use to compute and evaluate the CVaR of utilitarian and egalitarian welfare. For the logistic model, we build a polyhedral uncertainty set by estimating the cross-entropy loss on a held-out test set, and for the Gaussian process model we simply consider the confidence intervals of the resulting normal distribution.

For the binarized bids, we first set aside some of the observed bids as a test set. We estimate the missing bids and the bids for the held-out test pairs using logistic matrix factorization. Setting a hidden dimension size d, we construct two matrices $X \in \mathbb{R}^{n \times d}$ and $Y \in \mathbb{R}^{m \times d}$. We set $d = 20$. Let V^* denote the true binarized bid matrix, where we observe entries for the training set pairs $(a,i) \in \mathcal{T}$. We predict the probability of an affirmative bid as $\sigma((XY^\intercal)_{a,i})$ where σ is the logistic sigmoid function. We select X and Y to minimize the loss function

$$\sum_{(a,i) \in \mathcal{T}} -V^*_{a,i} \ln \left(\sigma((XY^\intercal)_{a,i}) \right) - V^*_{a,i} \ln \left((XY^\intercal)_{a,i} \right) \ .$$

For CVaR, we take samples from the distribution defined by $\sigma(XY^\intercal)$, assuming all pairs are independently-distributed. We also construct an uncertainty set as described in Appendix D using the cross-entropy loss on the test pairs.

Under the Gaussian process matrix factorization model [35], we simply predict a mean and variance of a normal distribution for each reviewer-paper pair. We can then sample values independently for each pair, or give a confidence interval for the joint Gaussian with $mn - 1$ degrees of freedom.

F Grouping Papers and Reviewers

We group papers and reviewers as follows: given the real-valued bids in the set $\{0.01, .5, 1\}$ we set unknown bids to be 0. We then construct a graph with all reviewers and papers as nodes, and the bid score between reviewers and papers is the edge weight. All inter-reviewer and inter-paper edges are set to 0 edge weight. We apply spectral embedding with 5 dimensions to transform the nodes into vectors, and cluster the resulting vectors into 4 clusters to obtain 4 groups containing both papers and reviewers. To ensure a balance of reviewers and papers across clusters, we employ Lloyd's algorithm for KMeans clustering with the modification that during each assignment step we enforce a lower bound on the number of papers and number of reviewers assigned to each cluster.

G Additional Experiments

For the binarized AAMAS 2016 and 2021 datasets, Tables 3 and 4 show the performance of the baseline USW and GESW maximizing allocations, the $\mathrm{CVaR}_0.01$ USW and GESW maximizing

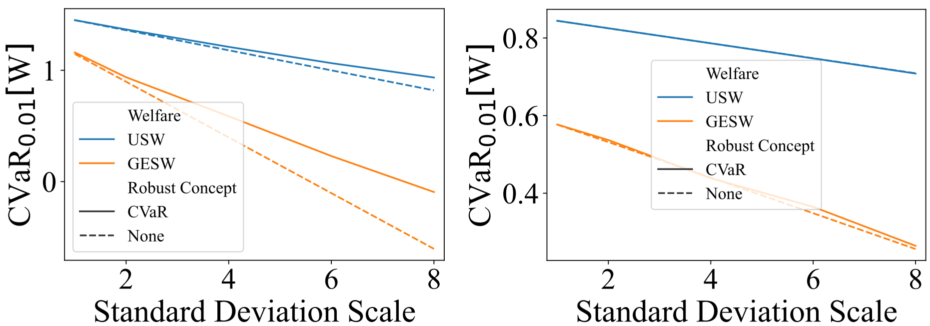

(a) AAMAS 2016 (b) AAMAS 2021

Figure 2: $\mathrm{CVaR}_{0.01}$ as noise increases for AAMAS 2016 and 2021.

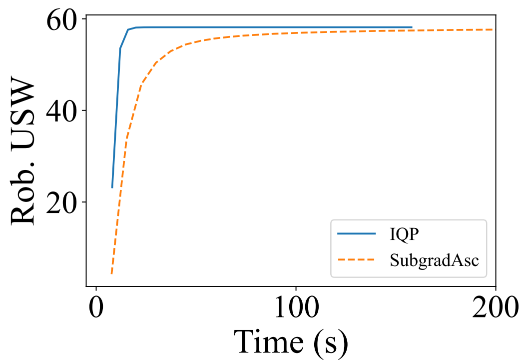

(a) AAMAS 2016

Figure 3: Convergence of the IQP vs. subgradient ascent on AAMAS 2016 dataset for the adversarial USW objective. The IQP (in blue) converges much faster.

allocations, and the adversarially-robust USW and GESW maximizing allocations at the $\delta = 0.3$ level. Because so many of the bids in AAMAS 2021 are recorded as `no`, since `no` is the default bid, we randomly select 90% of the `no` bids to be converted to `no response`.

Tables 5 to 7 show the same results for the Gaussian matrix factorization version of the 3 datasets, with the $\mathrm{CVaR}_{0.01}$ estimated by sampling from the estimated Gaussian distribution, and the adversarial welfare computed over the truncated ellipsoidal uncertainty set corresponding to the $1 - \delta$ confidence interval of the Gaussian. Results are not reported for the adversarial GESW approach, since the basic subgradient ascent approach fails to converge even after $1,000$ iterations.

Table 3: Performance of different allocations across each metric on the AAMAS 2016 dataset.

Allocation	Evaluation Objective					
	USW	GESW	CVaR USW	CVaR GESW	Rob. USW	Rob. GESW
USW	$\mathbf{1.00 \pm 0}$	$\mathbf{1.00 \pm 0}$	$\mathbf{1.00 \pm 0}$	$\mathbf{1.00 \pm 0}$	0 ± 0	0 ± 0
GESW	0.99 ± 0	$\mathbf{1.00 \pm 0}$	0.99 ± 0	0.99 ± 0.01	0 ± 0	0 ± 0
CVaR USW	0.99 ± 0	0.98 ± 0.01	0.99 ± 0	0.98 ± 0.01	0 ± 0	0 ± 0
CVaR GESW	0.99 ± 0.01	0.99 ± 0.01	0.98 ± 0.01	$\mathbf{1.00 \pm 0}$	0 ± 0	0 ± 0
Rob. USW	0.91 ± 0.02	0.87 ± 0.03	0.91 ± 0.02	0.90 ± 0.03	$\mathbf{1.00 \pm 0}$	$\mathbf{1.00 \pm 0}$
Rob. GESW	0.76 ± 0.05	0.66 ± 0.04	0.76 ± 0.05	0.65 ± 0.05	0.74 ± 0.10	$\mathbf{1.00 \pm 0}$

Table 4: Performance of different allocations across each metric on the AAMAS 2021 dataset.

Allocation	Evaluation Objective					
	USW	GESW	CVaR USW	CVaR GESW	Rob. USW	Rob. GESW
USW	$\mathbf{1.00 \pm 0}$	$\mathbf{1.00 \pm 0}$	$\mathbf{1.00 \pm 0}$	$\mathbf{1.00 \pm 0}$	0 ± 0	0.40 ± 0.49
GESW	$\mathbf{1.00 \pm 0}$	$\mathbf{1.00 \pm 0}$	$\mathbf{1.00 \pm 0}$	$\mathbf{1.00 \pm 0}$	0 ± 0	0.40 ± 0.49
CVaR USW	$\mathbf{1.00 \pm 0}$	$\mathbf{1.00 \pm 0}$	$\mathbf{1.00 \pm 0}$	$\mathbf{1.00 \pm 0}$	0 ± 0	0.40 ± 0.49
CVaR GESW	$\mathbf{1.00 \pm 0}$	$\mathbf{1.00 \pm 0}$	0.99 ± 0	$\mathbf{1.00 \pm 0}$	0 ± 0	0.40 ± 0.49
Rob. USW	0.85 ± 0.04	0.69 ± 0.14	0.84 ± 0.05	0.64 ± 0.19	$\mathbf{1.00 \pm 0}$	$\mathbf{1.00 \pm 0}$
Rob. GESW	0.48 ± 0.09	0.32 ± 0.12	0.43 ± 0.09	0.20 ± 0.12	0.07 ± 0.08	$\mathbf{1.00 \pm 0}$

Table 5: Performance of different allocations across each metric on the Gaussian AAMAS 2015 dataset.

Allocation	Evaluation Objective					
	USW	GESW	CVaR USW	CVaR GESW	Rob. USW	Rob. GESW
USW	$\mathbf{1.00 \pm 0}$	0.95 ± 0.03	$\mathbf{1.00 \pm 0}$	0.94 ± 0.04	0.61 ± 0.19	0.34 ± 0.34
GESW	0.87 ± 0.08	$\mathbf{1.00 \pm 0}$	0.86 ± 0.09	0.98 ± 0.02	0.42 ± 0.30	0.32 ± 0.35
CVaR USW	$\mathbf{1.00 \pm 0}$	0.94 ± 0.03	$\mathbf{1.00 \pm 0}$	0.96 ± 0.04	0.63 ± 0.19	0.36 ± 0.33
CVaR GESW	0.90 ± 0.06	0.99 ± 0.01	0.90 ± 0.07	$\mathbf{1.00 \pm 0}$	0.51 ± 0.26	0.36 ± 0.33
Rob. USW	0.86 ± 0.07	0.76 ± 0.12	0.88 ± 0.06	0.80 ± 0.10	$\mathbf{1.00 \pm 0}$	$\mathbf{1.00 \pm 0}$

Table 6: Performance of different allocations across each metric on the Gaussian AAMAS 2016 dataset.

Allocation	Evaluation Objective					
	USW	GESW	CVaR USW	CVaR GESW	Rob. USW	Rob. GESW
USW	$\mathbf{1.00 \pm 0}$	0.99 ± 0.01	$\mathbf{1.00 \pm 0}$	0.99 ± 0.02	0.47 ± 0.27	0.25 ± 0.38
GESW	0.91 ± 0.06	$\mathbf{1.00 \pm 0}$	0.91 ± 0.07	0.98 ± 0.01	0.37 ± 0.32	0.24 ± 0.38
CVaR USW	$\mathbf{1.00 \pm 0}$	0.98 ± 0.02	$\mathbf{1.00 \pm 0}$	0.99 ± 0.01	0.52 ± 0.25	0.27 ± 0.37
CVaR GESW	0.92 ± 0.05	0.98 ± 0.02	0.92 ± 0.06	$\mathbf{1.00 \pm 0}$	0.41 ± 0.31	0.28 ± 0.37
Rob. USW	0.84 ± 0.08	0.77 ± 0.12	0.86 ± 0.07	0.84 ± 0.09	$\mathbf{1.00 \pm 0}$	$\mathbf{1.00 \pm 0}$

Table 7: Performance of different allocations across each metric on the Gaussian AAMAS 2021 dataset.

Allocation	Evaluation Objective					
	USW	GESW	CVaR USW	CVaR GESW	Rob. USW	Rob. GESW
USW	$\mathbf{1.00 \pm 0}$	$\mathbf{1.00 \pm 0.01}$	$\mathbf{1.00 \pm 0}$	$\mathbf{1.00 \pm 0.01}$	0.53 ± 0.26	0.21 ± 0.40
GESW	0.80 ± 0.12	$\mathbf{1.00 \pm 0}$	0.79 ± 0.12	0.99 ± 0.01	0.24 ± 0.39	0.20 ± 0.40
CVaR USW	$\mathbf{1.00 \pm 0}$	$\mathbf{1.00 \pm 0.01}$	$\mathbf{1.00 \pm 0}$	$\mathbf{1.00 \pm 0.01}$	0.53 ± 0.26	0.21 ± 0.40
CVaR GESW	0.85 ± 0.08	$\mathbf{1.00 \pm 0}$	0.84 ± 0.08	$\mathbf{1.00 \pm 0}$	0.36 ± 0.34	0.20 ± 0.40
Rob. USW	0.81 ± 0.11	0.69 ± 0.16	0.81 ± 0.11	0.71 ± 0.16	$\mathbf{1.00 \pm 0}$	$\mathbf{1.00 \pm 0}$

H Proofs

H.1 Proof of Proposition 3.1

Proposition 3.1. *The problem in* (1) *is equivalent to solving*

$$\max_{\substack{\boldsymbol{\xi}\in\Lambda,\boldsymbol{\lambda}\in\mathbb{R}^l_{0+},\\ \boldsymbol{\beta}\in\mathbb{R}^k_{0+}}} \mathbf{c}^\intercal\Sigma_{[L]}^{-1}\mathbf{d}^\intercal + \boldsymbol{\beta}^\intercal\mathbf{e} - \frac{1}{4}\|\mathbf{c}^\intercal\Sigma_{[L]}^{-1/2}\|_2^2 + \sum_{i=1}^{l}\boldsymbol{\lambda}_i\|\bar{\mathbf{v}}_i^\intercal\Sigma_i^{-1/2}\|_2^2 - \|\mathbf{d}^\intercal\Sigma_{[L]}^{-1/2}\|_2^2 - \sum_{i=1}^{l}\boldsymbol{\lambda}_i r_i^2 \ , \tag{2}$$

where $\mathbf{c} = (-\boldsymbol{\beta}^\intercal\boldsymbol{Q}+\boldsymbol{\xi})$ *and* $\mathbf{d} = \sum_{i=1}^{l}\boldsymbol{\lambda}_i\bar{\mathbf{v}}_i^\intercal\Sigma_i^{-1}$*, and* $\Sigma_{[L]} = \sum_{i=1}^{l}\boldsymbol{\lambda}_i\Sigma_i^{-1}$*. Let* $\boldsymbol{\xi}^*$ *be the optimal* $\boldsymbol{\xi}$ *in* (2)*. Then, the optimal allocation* $\mathbf{a}^*$ *can be derived from* $\boldsymbol{\xi}^*$ *by solving the system of equations:*

$$G\in\mathcal{G}: \frac{\boldsymbol{w}_G}{|G|}\cdot\mathbf{a}_G \preceq \boldsymbol{\xi}^*_G,\ \mathbf{a}\in\mathcal{A} \quad ,$$

Proof. Let $\forall G\in\mathcal{G}, \mathbf{a}'_G = \frac{\boldsymbol{w}_G\cdot\mathbf{a}_G}{|G|}$. Consider the inner-minimization problem:

$$\begin{aligned} &\min_{\mathbf{v}\in\mathbb{R}^{nm}} \sum_{G\in\mathcal{G}}\boldsymbol{w}_G\cdot\mathbf{u}(\mathbf{a},\mathbf{v},G)\\ &\forall i\in[1,l], (\mathbf{v}-\bar{\mathbf{v}}_i)\Sigma_i^{-1/2}(\mathbf{v}-\bar{\mathbf{v}}_i)\le r_i^2\\ &\boldsymbol{Q}\mathbf{v}\succeq\mathbf{e}\\ &\mathbf{v}\succeq 0 \ , \end{aligned} \tag{14}$$

We will use the Lagrangian method for computing the dual of the above problem. The Lagrangian for the above problem is given by

$$\begin{aligned} L(\mathbf{v},\boldsymbol{\lambda}\in\mathbb{R}^l_{0+},\boldsymbol{\beta}\in\mathbb{R}^k,\boldsymbol{\zeta}\in\mathbb{R}^{nm}) = \mathbf{a}'^\intercal\mathbf{v} + \sum_{i=1}^{l}\boldsymbol{\lambda}_i\left(\mathbf{v}-\bar{\mathbf{v}}_i)\Sigma_i^{-1}(\mathbf{v}-\bar{\mathbf{v}}_i) - r_i^2\right)\\ - \boldsymbol{\beta}^\intercal(Q\mathbf{v}-\mathbf{e}) - \boldsymbol{\zeta}^\intercal\mathbf{v} \ . \end{aligned} \tag{15}$$

From the first-order optimality conditions, we get

$$\begin{aligned} \frac{\partial L(\mathbf{v},\boldsymbol{\lambda}\in\mathbb{R}^l_{0+},\boldsymbol{\beta}\in\mathbb{R}^k,\boldsymbol{\zeta}\in\mathbb{R}^{nm})}{\partial V} &= 0\\ \mathbf{a}' + \sum_{i=1}^{l}2\boldsymbol{\lambda}_i(\mathbf{v}-\bar{\mathbf{v}}_i)\Sigma_i^{-1} - \boldsymbol{\beta}^\intercal Q - \boldsymbol{\zeta} &= 0\\ \implies \mathbf{v} = \frac{\sum_{i=1}^{l}2\boldsymbol{\lambda}_i\bar{\mathbf{v}}_i^\intercal\Sigma_i^{-1} - (\mathbf{a}'-\boldsymbol{\beta}^\intercal Q-\boldsymbol{\zeta})}{\sum_{i=1}^{l}2\boldsymbol{\lambda}_i\Sigma_i^{-1}} \ . \end{aligned}$$

Substituting this value of $\mathbf{v}$ in (15), we get,

$$\begin{aligned} \max_{\substack{\boldsymbol{\lambda}\in\mathbb{R}^l,\\ \boldsymbol{\beta}\in\mathbb{R}^{nm},\\ \boldsymbol{\zeta}\in\mathbb{R}^{nm}}} &-\frac{1}{4}\left(\left(\mathbf{a}'-\boldsymbol{\beta}^\intercal Q-\boldsymbol{\zeta}\right)^\intercal\left(\sum_{i=1}^{l}\boldsymbol{\lambda}_i\Sigma_i^{-1}\right)^{-1}\left(\mathbf{a}'-\boldsymbol{\beta}^\intercal Q-\boldsymbol{\zeta}\right)\right) + \sum_{i=1}^{l}\boldsymbol{\lambda}_i\bar{\mathbf{v}}_i^\intercal\Sigma_i^{-1}\bar{\mathbf{v}}_i\\ &-\left(\sum_{i=1}^{l}\boldsymbol{\lambda}_i\bar{\mathbf{v}}_i\Sigma_i^{-1}\right)\left(\sum_{i=1}^{l}\boldsymbol{\lambda}_i\Sigma_i^{-1}\right)^{-1}\left(\sum_{i=1}^{l}\boldsymbol{\lambda}_i\bar{\mathbf{v}}_i\Sigma_i^{-1}\right)^\intercal\\ &+(\mathbf{a}'-\boldsymbol{\beta}^\intercal Q-\boldsymbol{\zeta})^\intercal\left(\sum_{i=1}^{l}\boldsymbol{\lambda}_i\Sigma_i^{-1}\right)^{-1}\left(\sum_{i=1}^{l}\boldsymbol{\lambda}_i\bar{\mathbf{v}}_i^\intercal\Sigma_i^{-1}\right)^\intercal - \sum_{i=1}^{l}\boldsymbol{\lambda}_i r_i^2 + \boldsymbol{\beta}^\intercal\mathbf{e}\\ &\boldsymbol{\lambda}\succeq 0\\ &\boldsymbol{\beta}\succeq 0\\ &\boldsymbol{\zeta}\ge 0 \end{aligned} \tag{16}$$

Using change of variables $\boldsymbol{\zeta} = \mathbf{a}' - \boldsymbol{\xi}$, and combining the dual with the outer-maximization problem in (1), we get

$$\begin{aligned}
\max_{\substack{\mathbf{a}\in\mathcal{A},\boldsymbol{\lambda}\in\mathbb{R}^l_{0+},\boldsymbol{\beta}\in\mathbb{R}^k_{0+},\\ \boldsymbol{\zeta}\in\mathbb{R}^{nm}_{0+},\boldsymbol{\xi}\in\mathbb{R}^{nm}}} \quad & -\frac{1}{4}\|\mathbf{c}^\intercal\Sigma_{[L]}^{-1/2}\|_2^2 + \sum_{i=1}^{l}\boldsymbol{\lambda}_i\|\bar{\mathbf{v}}_i^\intercal\Sigma_i^{-1/2}\|_2^2 \\
& - \|\mathbf{d}^\intercal\Sigma_{[L]}^{-1/2}\|_2^2 + \mathbf{c}^\intercal\Sigma_{[L]}^{-1}\mathbf{d}^\intercal - \sum_{i=1}^{l}\boldsymbol{\lambda}_i r_i^2 + \boldsymbol{\beta}^\intercal\mathbf{e} \\
\text{s.t. } & \boldsymbol{\zeta} = \mathbf{a}' - \boldsymbol{\xi} \ ,
\end{aligned} \tag{17}$$

where $\mathbf{c} = (\boldsymbol{\beta}^\intercal Q + \boldsymbol{\xi})$, $\Sigma_{[L]} = \left(\sum_{i=1}^{l}\boldsymbol{\lambda}_i\Sigma_i^{-1}\right)$, and $\mathbf{d} = \sum_{i=1}^{l}\boldsymbol{\lambda}_i\bar{\mathbf{v}}_i^\intercal\Sigma_i^{-1}$. Note that the above optimization problem is concave; from affine-composition rule in convex optimization, we retain the concavity of the objective after the change of variable and the allocation $\mathbf{a}$ only appears in a linear constraint which is convex.

We further simplify the above problem by eliminating the allocation variables $\mathbf{a}$ and the dual variable $\boldsymbol{\zeta}$ and subsequently deriving them from the solution of the resultant problem.

Note that in the above problem $\mathbf{a}' - \boldsymbol{\zeta} = \boldsymbol{\xi}$. Let $(\mathbf{a}^*, \boldsymbol{\zeta}^*)$ represent an optimal $(\mathbf{a}, \boldsymbol{\zeta})$ pair for the problem in (17). Note that there can be multiple pairs of $(\mathbf{a}, \boldsymbol{\zeta})$ that are optimal. Let $\mathbf{a}'^*_G = \frac{\boldsymbol{w}_G}{|G|}\mathbf{a}^*_G, \forall G \in \mathcal{G}$. To eliminate $\boldsymbol{\zeta}$ and $\mathbf{a}$, we need to find a set of feasible $\boldsymbol{\xi}$, which we denote by Λ, such that there existing a $\boldsymbol{\xi}' \in \Xi$ such that $\boldsymbol{\xi}' = \mathbf{a}'^* - \boldsymbol{\zeta}^*$ for at least one optimal pair $(\mathbf{a}^*, \boldsymbol{\zeta}^*)$. It is easy to see that if there exists such a $\boldsymbol{\xi}' \in \Xi$, then, $\boldsymbol{\xi}'$ maximizes the objective in (17). Furthermore, it is easy to verify that $\Lambda = \mathcal{A} = \mathbb{R}^{nm}_{0+} = \{\boldsymbol{\xi} \in \mathbb{R}^{nm} \mid \forall a \in N : \sum_{i\in I}\boldsymbol{\xi}_{am+i} \leq \bar{\kappa}_a, \forall i \in I : \sum_{a\in N}\boldsymbol{\xi}_{am+i} \leq \bar{\psi}_i, \boldsymbol{\xi} \preceq \boldsymbol{c}\}$ satisfies the above criteria for optimality.

Thus, we can break down the problem in (17) into two sub-problems. In the first problem, we obtain the optimal value of $\boldsymbol{\lambda}$, $\boldsymbol{\xi}$ and $\boldsymbol{\beta}$ by solving:

$$\begin{aligned}
\boldsymbol{\zeta}^*, \boldsymbol{\beta}^*, \boldsymbol{\xi}^* = \operatorname*{arg\,max}_{\substack{\boldsymbol{\zeta}\in\mathbb{R}^{nm}_{0+},\boldsymbol{\beta}\in\mathbb{R}^k_{0+},\\ \boldsymbol{\xi}\in\Lambda}} \quad & -\frac{1}{4}\|\mathbf{c}^\intercal\Sigma_{[L]}^{-1/2}\|_2^2 + \sum_{i=1}^{l}\boldsymbol{\lambda}_i\|\bar{\mathbf{v}}_i^\intercal\Sigma_i^{-1/2}\|_2^2 - \|\mathbf{d}^\intercal\Sigma_{[L]}^{-1/2}\|_2^2 \\
& + \mathbf{c}^\intercal\Sigma_{[L]}^{-1}\mathbf{d}^\intercal - \sum_{i=1}^{l}\boldsymbol{\lambda}_i r_i^2 + \boldsymbol{\beta}^\intercal\mathbf{e} \ ,
\end{aligned}$$

where $\mathbf{c} = (-\boldsymbol{\beta}^\intercal\boldsymbol{Q} + \boldsymbol{\xi})$ and $\mathbf{d} = \sum_{i=1}^{l}\boldsymbol{\lambda}_i\bar{\mathbf{v}}_i^\intercal\Sigma_i^{-1}$, and $\Sigma_{[L]} = \left(\sum_{i=1}^{l}\boldsymbol{\lambda}_i\Sigma_i^{-1}\right)$. Then, the set of optimal $(\mathbf{a}, \boldsymbol{\zeta})$ pairs are computed by solving a system of equations: $\{(\mathbf{a}, \boldsymbol{\zeta}) \mid \mathbf{a} \in \mathcal{A}, G \in \mathcal{G}, \frac{\boldsymbol{w}_G}{|G|} \cdot \mathbf{a}_G - \boldsymbol{\zeta}_G = \boldsymbol{\xi}^*_G, \boldsymbol{\zeta} \in \mathbb{R}^{nm}_{0+}\}$. □

H.2 Proof of proposition 3.2

Proposition 3.2. *In the case where the uncertainty set $\mathcal{V}$ is defined purely by linear constraints, i.e., $\mathcal{V} = \{\mathbf{v} \in \mathbb{R}^{nm}_{0+} \mid \boldsymbol{Q}\mathbf{v} \succeq \mathbf{e}\}$, the optimal allocation $\mathbf{a}^*$ for the problem in* (1) *can be computed by solving the linear program:*

$$\max_{\mathbf{a}\in\mathcal{A},\boldsymbol{\beta}\in\mathbb{R}^k_{0+}} \boldsymbol{\beta}^\intercal\mathbf{e} \quad \textit{s.t. } \forall G \in \mathcal{G} : \ \boldsymbol{\beta}_G^\intercal\boldsymbol{Q}_G \preceq \frac{\boldsymbol{w}_G}{|G|}\mathbf{a}_G \ . \tag{3}$$

Proof. Consider the following inner-minimization problem. Let $\forall G \in \mathcal{G}, \mathbf{a}'_G = \frac{\boldsymbol{w}_G \cdot \mathbf{a}_G}{|G|}$. Consider the inner-minimization problem:

$$\begin{aligned}
\min_{\mathbf{v}\in\mathbb{R}^{nm}} \ & \sum_{G\in\mathcal{G}} \boldsymbol{w}_G \cdot \mathbf{u}(\mathbf{a}, \mathbf{v}, G) \\
& \boldsymbol{Q}\mathbf{v} \succeq \mathbf{e} \\
& \mathbf{v} \succeq 0 \ ,
\end{aligned}$$

We compute the dual of the above problem using the Lagrangian method.

$$
\begin{aligned}
L(\mathbf{v}, \boldsymbol{\lambda} \in \mathbb{R}^{l}_{0+}, \boldsymbol{\beta} \in \mathbb{R}^{k}_{0+}, \boldsymbol{\zeta} \in \mathbb{R}^{nm}_{0+}) &= \mathbf{a}'^{\mathsf{T}}\mathbf{v} - \boldsymbol{\beta}^{\mathsf{T}}(\boldsymbol{Q}\mathbf{v} - \mathbf{e}) - \boldsymbol{\zeta}^{\mathsf{T}}\mathbf{v} \\
&= (\mathbf{a}' - \boldsymbol{\beta}^{\mathsf{T}}\boldsymbol{Q} - \boldsymbol{\zeta})^{\mathsf{T}}\mathbf{v} + \boldsymbol{\beta}^{\mathsf{T}}\mathbf{e} \\
L(\boldsymbol{\lambda} \in \mathbb{R}^{l}_{0+}, \boldsymbol{\beta} \in \mathbb{R}^{k}_{0+}, \boldsymbol{\zeta} \in \mathbb{R}^{nm}_{0+}) &= \begin{cases} \boldsymbol{\beta}^{\mathsf{T}}\mathbf{e} & (\mathbf{a}' - \boldsymbol{\beta}^{\mathsf{T}}\boldsymbol{Q} - \boldsymbol{\zeta}) \succeq 0 \\ -\infty & \text{otherwise} \end{cases}
\end{aligned} \tag{18}
$$

Therefore, the dual is given by

$$
\begin{aligned}
\max_{\boldsymbol{\beta} \in \mathbb{R}^{k}_{0+}, \boldsymbol{\zeta} \in \mathbb{R}^{nm}_{0+}} \quad & \boldsymbol{\beta}^{\mathsf{T}}\mathbf{e} \\
& \boldsymbol{\beta}^{\mathsf{T}}\boldsymbol{Q} - \boldsymbol{\zeta} \preceq \mathbf{a}' \ .
\end{aligned} \tag{19}
$$

Since $\boldsymbol{\zeta}$ is non-negative, we can eliminate it to get

$$
\begin{aligned}
\max_{\boldsymbol{\beta} \in \mathbb{R}^{k}_{0+}} \quad & \boldsymbol{\beta}^{\mathsf{T}}\mathbf{e} \\
& \boldsymbol{\beta}^{\mathsf{T}}Q \preceq \mathbf{a}' \ .
\end{aligned} \tag{20}
$$

Combining the dual with the outer-maximization problem in (1) yields the final result.

□

H.3 Proof of proposition 3.3

Proposition 3.3. *Suppose that the set $\mathcal{V}$ in* (1) *is defined by a single truncated ellipsoidal constraint i.e., $\mathcal{V} = \{\mathbf{v} \in \mathbb{R}^{nm}_{0+} \mid (\mathbf{v} - \bar{\mathbf{v}})\Sigma_i^{-1}(\mathbf{v} - \bar{\mathbf{v}}) \leq r^2\}$. The problem in* (1) *is equivalent to solving*

$$
\max_{\lambda \in \mathbb{R}_{0+}, \boldsymbol{\xi} \in \Lambda} \left(\boldsymbol{\xi}^{\mathsf{T}}\bar{\mathbf{v}} - \frac{\|\boldsymbol{\xi}^{\mathsf{T}}\Sigma^{\frac{1}{2}}\|_F^2}{4\lambda} - \lambda r^2 \right) \ . \tag{4}
$$

The exact optimal solution $(\lambda^*, \boldsymbol{\xi}^*)$ to eq. (4) *can be computed by alternately performing two steps until convergence: first, fixing $\boldsymbol{\xi}$ and optimizing λ, i.e., $\lambda = \|\boldsymbol{\xi}^{\mathsf{T}}\Sigma^{\frac{1}{2}}\|_F^2/2r$, and second, fixing λ and solving a concave quadratic program to optimize $\boldsymbol{\xi}$. Furthermore, the optimal allocation $\mathbf{a}^*$ can be computed from $\boldsymbol{\xi}^*$ as in Proposition 3.1.*

Proof. Setting $k = 0$ and $l = 1$ in (2) yields the stated optimization problem. In appendix H.4, we established that this dual is concave in λ and $\boldsymbol{\xi}$. Unfortunately, the objective in (4) is a cubic polynomial that is difficult to optimize exactly using standard solvers. However, since the objective is concave and differentiable, we can leverage a block coordinate descent like technique to achieve the global optimal solution, i.e., we can alternate between optimizing $(\boldsymbol{\zeta}, \boldsymbol{\xi}, \mathbf{a})$, which is a quadratic problem, and optimizing $(\boldsymbol{\lambda})$ which has a closed form solution $\lambda = \frac{\|\boldsymbol{\xi}^{\mathsf{T}}\Sigma^{\frac{1}{2}}\|_F^2}{2r}$, until convergence. Since the objective is concave and differentiable, the above algorithm is guaranteed to converge to the global optimal solution [11].

□

H.4 Proof of proposition 3.4

Proposition 3.4. *The problem in* (5) *is equivalent to solving*

$$
\max_{\substack{\boldsymbol{\xi} \in \Lambda, \\ \boldsymbol{\lambda} \in \mathbb{R}^{g \times l}_{0+}, \\ \boldsymbol{\beta} \in \mathbb{R}^{g \times k}_{0+}}} \min_{G \in \mathcal{G}} \boldsymbol{\beta}_G^{\mathsf{T}}\mathbf{e}_G + \mathbf{c}_G^{\mathsf{T}}\Sigma_*^{-1}\mathbf{d}_G - \frac{1}{4}\|\mathbf{c}_G^{\mathsf{T}}\Sigma_*^{-1/2}\|_2^2 + \sum_{i=1}^{l}\left(\boldsymbol{\lambda}_{G,i}\|\bar{\mathbf{v}}_{G,i}^{\mathsf{T}}\Sigma_*^{-1/2}\|_2^2 - \boldsymbol{\lambda}_{G,i}r_{G,i}^2\right) - \|\mathbf{d}_G^{\mathsf{T}}\Sigma_*^{-1/2}\|_2^2 \tag{7}
$$

and $\forall G \in \mathcal{G}:\ \mathbf{c}_G = (\boldsymbol{\xi}_G - \boldsymbol{\beta}_G^{\mathsf{T}}\boldsymbol{Q}_G)$, $\mathbf{d}_G = \sum_{i=1}^{l}\boldsymbol{\lambda}_{G,i}\bar{\mathbf{v}}_{G,i}\Sigma_{G,i}^{-1}$, and $\Sigma_* = \sum_{i=1}^{l},\ \boldsymbol{\lambda}_{G,i}\Sigma_{G,i}^{-1}$. Let $\boldsymbol{\xi}^*$ be the optimal $\boldsymbol{\xi}$ in (7)*. Then, the optimal allocation $\mathbf{a}^*$ satisfies the system of equations:*

$$
G \in \mathcal{G}: \frac{\boldsymbol{w}_G}{|G|} \cdot \mathbf{a}_G \preceq \boldsymbol{\xi}_G^*,\ \mathbf{a} \in \mathcal{A} \quad ,
$$

Proof. Consider the following optimization problem.

$$\begin{aligned}\max_{\mathbf{a}\in\mathcal{A}}\ \min_{G\in\mathcal{G}}\ \min_{\mathbf{v}_G\in\mathcal{V}_G}\ & \frac{1}{|G|}\mathbf{a}_G^{\intercal}\mathbf{v}_G \\ & \forall i\in[1,l],\ \forall G\in\mathcal{G}\ (\mathbf{v}_G-\bar{\mathbf{v}}_{i,G})\Sigma_{i,G}^{-1/2}(\mathbf{v}_G-\bar{\mathbf{v}}_{i,G})\leq r_{i,G}^2 \\ & \forall G\in\mathcal{G},\ Q_G\mathbf{v}_G\succeq\mathbf{e}_G \\ & \mathbf{v}_G\succeq 0\ \ .\end{aligned}\tag{21}$$

It is important to note that the inner-minimization is a convex optimization problem and the outer-maximization is a concave maximization problem. This is due to the fact that affine functions are either concave or convex and minimum of concave objectives is concave.

Furthermore, the inner-most minimization over the uncertainty set of valuation matrices is independent for each group. Thus, simply replacing each of these minimization problems with their respective duals yields the following problem.

$$\begin{aligned}\max_{\mathbf{a}\in\mathcal{A}}\ \min_{G\in\mathcal{G}}\ \max_{\substack{\boldsymbol{\lambda}_G\in\mathbb{R}^l,\\ \boldsymbol{\beta}_G\in\mathbb{R}^{nm},\\ \boldsymbol{\zeta}_G\in\mathbb{R}^{nm}}}\ & -\frac{1}{4}\left(\left(\mathbf{a}'_G-\boldsymbol{\beta}_G^{\intercal}\boldsymbol{Q}_G-\boldsymbol{\zeta}_G\right)^{\intercal}(\sum_{i=1}^{l}\lambda_i\Sigma_i^{-1})^{-1}\left(\mathbf{a}'_G-\boldsymbol{\beta}_G^{\intercal}Q_G-\boldsymbol{\zeta}_G\right)\right)+ \\ & \sum_{i=1}^{l}\lambda_{G,i}\bar{\mathbf{v}}_{G,i}^{\intercal}\Sigma_{G,i}^{-1}\bar{\mathbf{v}}_{G,i}-\left(\sum_{i=1}^{l}\lambda_{G,i}\bar{\mathbf{v}}_{G,i}\Sigma_{G,i}^{-1}\right)\left(\sum_{i=1}^{l}\lambda_{G,i}\Sigma_{G,i}^{-1}\right)^{-1}\left(\sum_{i=1}^{l}\lambda_{G,i}\bar{\mathbf{v}}_{G,i}\Sigma_{G,i}^{-1}\right)^{\intercal} \\ & +(\mathbf{a}'_G-\boldsymbol{\beta}_G^{\intercal}\boldsymbol{Q}_G-\boldsymbol{\zeta}_G)^{\intercal}\left(\sum_{i=1}^{l}\lambda_{G,i}\Sigma_{G,i}^{-1}\right)^{-1}\left(\sum_{i=1}^{l}\lambda_{G,i}\bar{\mathbf{v}}_{G,i}^{\intercal}\Sigma_{G,i}^{-1}\right)^{\intercal} \\ & -\sum_{i=1}^{l}\lambda_{G,i}r_{G,i}^2+\boldsymbol{\beta}_G^{\intercal}\mathbf{e}_G \\ & \lambda_G\succeq 0 \\ & \boldsymbol{\beta}_G\succeq 0 \\ & \boldsymbol{\zeta}_G\geq 0\ \ .\end{aligned}\tag{22}$$

Using the change of variables $\boldsymbol{\zeta}_G=\mathbf{a}'_G-\boldsymbol{\xi}_G\forall G\in\mathcal{G}$, and combining the dual with the outer-maximization problem, we get

$$\begin{aligned}\max_{\mathbf{a}\in\mathcal{A}}\ \min_{G\in\mathcal{G}}\ \max_{\substack{\boldsymbol{\lambda}_G\in\mathbb{R}^l,\\ \boldsymbol{\beta}_G\in\mathbb{R}^{nm},\\ \boldsymbol{\zeta}_G\in\mathbb{R}^{nm}}}\ & -\frac{1}{4}\|\mathbf{c}_G^{\intercal}\Sigma_*^{-1/2}\|_2^2+\sum_{i=1}^{l}\boldsymbol{\lambda}_{G,i}\|\bar{\mathbf{v}}_{G,i}^{\intercal}\Sigma_{G,i}^{-1/2}\|_2^2 \\ & -\|\mathbf{d}_G^{\intercal}\Sigma_*^{-1/2}\|_2^2-\mathbf{c}_G^{\intercal}\Sigma_*^{-1}\mathbf{d}_G^{\intercal}-\sum_{i=1}^{l}\lambda_{G,i}r_{G,i}^2+\boldsymbol{\beta}_G^{\intercal}\mathbf{e}_G \\ & \text{s.t. }\boldsymbol{\zeta}_G=\mathbf{a}'_G-\boldsymbol{\xi}_G\ \ ,\end{aligned}\tag{23}$$

where $\forall G\in\mathcal{G},\mathbf{c}_G=(\boldsymbol{\beta}_G^{\intercal}Q_G+\boldsymbol{\xi}_G)$, $\Sigma_*=\left(\sum_{i=1}^{l}\lambda_{G,i}\Sigma_{G,i}^{-1}\right)$, and $\mathbf{d}_G=\sum_{i=1}^{l}\lambda_{G,i}\bar{\mathbf{v}}_{G,i}^{\intercal}\Sigma_{G,i}^{-1}$.

Since the inner maximization for each group is independent of the other groups, we can re-order the inner minimization over groups and the inner-maximization problem. Thus, without loss of generality, we can write the above optimization problem as

$$\begin{aligned}\max_{\substack{\mathbf{a}\in\mathcal{A},\boldsymbol{\zeta}\in\mathbb{R}^{nm},\boldsymbol{\lambda}\in\mathbb{R}_{0+}^{g\times l},\\ \boldsymbol{\beta}\in\mathbb{R}_{0+}^{g\times k},\boldsymbol{\xi}\in\mathbb{R}^{nm}}}\ \min_{G\in\mathcal{G}}\ & -\frac{1}{4}\|\mathbf{c}_G^{\intercal}\Sigma_*^{-1/2}\|_2^2+\sum_{i=1}^{l}\boldsymbol{\lambda}_{i,G}\|\bar{\mathbf{v}}_{i,G}^{\intercal}\Sigma_*^{-1/2}\|_2^2-\|\mathbf{d}_G^{\intercal}\Sigma_*^{-1/2}\|_2^2 \\ & -\mathbf{c}_G^{\intercal}\Sigma_*^{-1}\mathbf{d}_G^{\intercal}-\sum_{i=1}^{l}\boldsymbol{\lambda}_{G,i}r_{i,G}^2+\boldsymbol{\beta}_G^{\intercal}\mathbf{e}_G \\ & \text{s.t. }\boldsymbol{\xi}_G=\frac{\mathbf{a}_G}{|G|}-\boldsymbol{\zeta}_G\ \ ,\end{aligned}\tag{24}$$

Using the same technique as in appendix H.1, we can simplify the problem by eliminating the variables $\mathbf{a}$ and $\boldsymbol{\zeta}$ in the above problem and later deriving them from the optimal $\boldsymbol{\xi}$.

Eliminating $\boldsymbol{\zeta}$ and $\mathbf{a}$ in (24), we get the following optimization problem.

$$\begin{aligned}\boldsymbol{\lambda}^*, \boldsymbol{\beta}^*, \boldsymbol{\xi}^* = \underset{,\boldsymbol{\lambda}\in\mathbb{R}_{0+}^{g\times l}, \boldsymbol{\beta}\in\mathbb{R}_{0+}^{g\times k}, \boldsymbol{\xi}\in\Lambda}{\arg\max} \min_{G\in\mathcal{G}} -\frac{1}{4}\|\mathbf{c}_G^\intercal \Sigma_*^{-1/2}\|_2^2 + \sum_{i=1}^{l} \boldsymbol{\lambda}_{G,i} \|\bar{\mathbf{v}}_{G,i}^\intercal \Sigma_*^{-1/2}\|_2^2 - \|\mathbf{d}_G^\intercal \Sigma_*^{-1/2}\|_2^2 \\ - \mathbf{c}_G^\intercal \Sigma_*^{-1} \mathbf{d}_G - \sum_{i=1}^{l} \boldsymbol{\lambda}_{G,i} r_{G,i}^2 + \boldsymbol{\beta}_G^\intercal \mathbf{e}_G \quad ,\end{aligned} \tag{25}$$

where $\forall G \in \mathcal{G},\ \mathbf{c}_G = (\boldsymbol{\xi}_G - \boldsymbol{\beta}_G^\intercal Q_G)$, $\mathbf{d}_G = \sum_{i=1}^{l} \lambda_{G,i} \bar{\mathbf{v}}_{G,i} \Sigma_{G,i}^{-1}$, and $\Sigma_* = \Big(\sum_{i=1}^{l},\ \lambda_{G,i}\Sigma_{G,i}^{-1}\Big)$.

□

H.5 Proof of proposition 3.5

Proposition 3.5. *In the case where the uncertainty set $\mathcal{V}$ is defined only by linear constraints, i.e., $\mathcal{V} = \{\mathbf{v} \in \mathbb{R}^{nm} \mid \boldsymbol{Q}\mathbf{v} \succeq \mathbf{e}, \mathbf{v} \succeq 0\}$, the max-min-min problem in* (5) *is trivially transformable into a linear program.*

Proof. Substituting $l = 0$ in (7), we get

$$\begin{aligned} \max_{\mathbf{a}\in\mathcal{A}, \boldsymbol{\beta}\in\mathbb{R}^{g\times k}} \min_{G\in\mathcal{G}} \quad & \boldsymbol{\beta}_G^\intercal \mathbf{e}_G \\ & \boldsymbol{\beta}_G^\intercal Q_G \preceq \mathbf{a}'_G \quad , \end{aligned} \tag{26}$$

where $\mathbf{a}' = \frac{\mathbf{a}}{|G|} \forall G \in \mathcal{G}$. Using simple algebraic manipulations, we can write the above optimization problem as

$$\begin{aligned} \max_{\mathbf{a}\in\mathcal{A}, \boldsymbol{\beta}\in\mathbb{R}^{g\times k}, t\in\mathbb{R}} \quad & t \\ \forall G \in \mathcal{G} : t \le \boldsymbol{\beta}_G^\intercal \mathbf{e}_G & \\ \forall G \in \mathcal{G} : \boldsymbol{\beta}_G^\intercal Q_G \preceq \mathbf{a}'_G & \quad , \end{aligned} \tag{27}$$

□

H.6 Proof of Proposition 3.6

Proposition 3.6. *Consider an optimization problem of the form*

$$\max_{\mathbf{a}\in\mathcal{A}} \min_{\mathbf{v}\in\mathcal{V}} W_M(\mathbf{u}(\mathbf{a}, \mathbf{v}, G_1), \mathbf{u}(\mathbf{a}, \mathbf{v}, G_2), \ldots, \mathbf{u}(\mathbf{a}, \mathbf{v}, G_g)) \quad , \tag{8}$$

where the welfare function $W_M(\cdot)$ is monotonic in the utility of groups. If the valuation uncertainty sets are independent across groups, $\mathcal{V} = \bigotimes_{G\in\mathcal{G}} \mathcal{V}_G$, then, the problem in (8) *simplifies to*

$$\max_{\mathbf{a}\in\mathcal{A}} W_M(\min_{\mathbf{v}\in\mathcal{V}} \mathbf{u}(\mathbf{a}, \mathbf{v}, G_1), \min_{\mathbf{v}\in\mathcal{V}} \mathbf{u}(\mathbf{a}, \mathbf{v}, G_2), \ldots, \min_{\mathbf{v}\in\mathcal{V}} \mathbf{u}(\mathbf{a}, \mathbf{v}, G_g)) \ .$$

Proof. The result directly follows from the monotonic property of the welfare function and the independence of the uncertainty sets across groups. □

H.7 Proof of proposition 4.1

Proposition 4.1. *Given h samples of $\tilde{\mathbf{v}}$, i.e., $\mathbf{v}^1, \mathbf{v}^2, \mathbf{v}^3, \ldots \mathbf{v}^h$ from $\mathcal{D}_{\mathbf{v}}$, the optimal allocation for the problem in* (9) *can be approximately computed by solving*

$$\max_{\mathbf{a}\in\mathcal{A}} \max_{\boldsymbol{y}\in\mathbb{R}_{0+}^h, b\in\mathbb{R}} \left(b - \frac{1}{\alpha}\sum_{j=1}^{h} \boldsymbol{y}_j\right) \qquad \forall j \in [1, h] :\ \boldsymbol{y}_j \ge \frac{1}{h}\left(b - \sum_{G\in\mathcal{G}} \frac{\boldsymbol{w}_G}{|G|} \cdot \mathbf{a}_G^\intercal \mathbf{v}_G^j\right) \quad . \tag{10}$$

Proof. For any random utility X, $\mathrm{CVaR}_\alpha[X]$ can be written as

$$\mathrm{CVaR}_\alpha[X] = \max_{b\in\mathbb{R}} b - \frac{1}{\alpha}\,\mathbb{E}\,[b-X]_+ \quad , \tag{28}$$

where $(t)_+ = max(t,0)$. Given a posterior distribution of valuations $\mathcal{D}_{\mathbf{v}}$, we generate h samples of the valuation matrix, i.e., $\mathbf{v}_1, \mathbf{v}_2 \ldots \mathbf{v}_h$, and use it to empirically estimate the expectation in (28).

$$\begin{aligned}
&\max_{\mathbf{a}\in\mathbb{R}^{n\times m}} \max_{y\in\mathbb{R}^h, b\in\mathbb{R}} \left(b - \frac{1}{\alpha}\sum_{j=1}^{h} y_j\right) \\
&\forall j \in [1,h]\; y_j \geq 0 \\
&\forall j \in [1,h]\; y_j \geq \frac{1}{h}b - \frac{1}{h}\sum_{G\in\mathcal{G}} \langle \mathbf{a}_G, \mathbf{v}_G^j\rangle_{\mathrm{F}} - \eta \sum_{G\in\mathcal{G}} \langle \mathbf{a}_G, \bar{\mathbf{v}}_G\rangle_{\mathrm{F}} \\
&\mathbf{a} \in \mathcal{A} \;\; .
\end{aligned} \tag{29}$$

□

H.8 Proof of proposition 4.3

Proposition 4.3. *Suppose that* $\mathbf{v}$ *is a multivariate sub-Gaussian with mean* $\bar{\mathbf{v}} \in \mathbb{R}^{nm}$ *and covariance proxy* $\Sigma \in \mathbb{R}^{nm\times nm}$*, i.e.,* $\exists K \geq 0$ *s.t.* $\mathbb{E}\left[\exp(\lambda(\mathbf{v}-\bar{\mathbf{v}})^\intercal \boldsymbol{z}))\right] \leq \exp(\lambda^2 K^2 \boldsymbol{z}^\intercal \Sigma \boldsymbol{z}/2), \forall \lambda \in \mathbb{R}, \forall \boldsymbol{z} \in \mathbb{R}^{nm}$ *and that Assumption 4.2 holds. Let* $|\mathcal{A}|$ *represent the number of feasible allocations,* $\mathbf{a}'_G = \frac{\boldsymbol{w}_G}{|G|}\cdot \mathbf{a}_G$*, and* h=$\mathcal{O}\left(\frac{8\max(\max_{\mathbf{a}\in\mathcal{A}} \mathbf{a}'^\intercal \Sigma \mathbf{a}', 8)\log\left(\frac{6|\mathcal{A}|}{\delta}\right)}{\varepsilon^2(\alpha)^2 \min(\eta^2,1)}\right)$ *where* $\delta \in (0,1)$*. Then,* $\Pr[\forall \mathbf{a}\in\mathcal{A} : |\hat{c}_{h,\alpha}(\mathbf{a}) - c_\alpha(\mathbf{a})| \leq \varepsilon] \geq 1-\delta$*.*

Proof.

Assumption H.1 (L.A. *et al.* [34])**.** *The random variable* X *is continuous with probability density function* f *that satisfies the following condition: There exists universal constants* $\eta, \delta' \geq 0$ *such that* $f(x) \geq \eta \; \forall x \in [v_\alpha - \frac{\delta'}{2}, v_\alpha + \frac{\delta'}{2}]$*, where* $v_\alpha = F^{-1}(\alpha)$*.*

Theorem H.2 (L.A. *et al.* [34])**.** *Let* $(X_i)_{i=1}^n$ *be a sequence of* $i.i.d$ *random variables. Let* $\hat{c}_{n,\alpha}$ *be the empirical CVaR estimates of X computed from the above samples. Suppose that* $X_i, i = 1, \ldots n$ *are* $\sigma-$*sub-Gaussian. Then for any* $\varepsilon \geq 0$*, we have*

$$\Pr\left[|c_{n,\alpha} - c_\alpha| > \varepsilon\right] \leq 6\exp\left(\frac{-n(\alpha)^2\min(\eta^2,1)}{8\max(8,\sigma^2)}\right) \quad . \tag{30}$$

From the assumption, we know that the valuation vector $\mathbf{v}$ is a sub-Gaussian that satisfies the following condition: $\exists K \geq 0$ s.t. $\mathbb{E}\left[\exp(\lambda(\mathbf{v}-\bar{\mathbf{v}})^\intercal \boldsymbol{w}))\right] \leq \exp(\lambda^2 K^2 \bar{\mathbf{v}}^\intercal \Sigma \bar{\mathbf{v}}/2), \forall \lambda \in \mathbb{R}, \forall \boldsymbol{w} \in \mathbb{R}^{nm}$.

Further, we know that if X is a σ-sub-Gaussian random variable then cX is also sub-Gaussian with variance proxy= $c\sigma$.

Using the above two properties, we get that the utilitarian welfare for a given allocation $\mathbf{a}$ is also a sub-Gaussian with variance-proxy = $K^2 \mathbf{a}'^\intercal \Sigma \mathbf{a}'$ where $\forall G \in \mathcal{G}, \mathbf{a}'_G = \frac{\boldsymbol{w}_G \mathbf{a}}{|G|}$.

For any allocation $\mathbf{a}$, let $\hat{c}_{h,\alpha}(\mathbf{a})$ represent the empirical estimate of CVaR of utilitarian welfare and $c_{h,\alpha}(\mathbf{a})$ represent the corresponding true value.

Then, we can bound the error of approximating the CVaR of the utilitarian welfare for allocation $\mathbf{a}$ as

$$\Pr\left[|\hat{c}_{M,\alpha}(\mathbf{a}) - c_\alpha(\mathbf{a})| > \varepsilon\right] \leq 6\exp\left(\frac{-h(\alpha)^2 min(\eta^2,1)}{8\max(8, K^2\mathbf{a}'^\intercal\Sigma\mathbf{a}')}\right) \tag{31}$$

Furthermore, the approximation error for all allocations can be upper-bounded as

$$\Pr\left[\forall \mathbf{a}\in\mathcal{A}, |\hat{c}_{h,\alpha}(\mathbf{a}) - c_\alpha(\mathbf{a})| \leq \varepsilon\right] \leq 1 - \sum_{\mathbf{a}\in\mathcal{A}} \Pr\left[|\hat{c}_{h,\alpha}(\mathbf{a}) - c_\alpha(\mathbf{a})| > \varepsilon\right] \tag{32}$$

Combining (31) and (32) and setting $h = \left(\frac{8\max(\max_{\mathbf{a}\in\mathcal{A}} \mathbf{a}'^{\intercal}\Sigma\mathbf{a}', 8)\log\left(\frac{6|\mathcal{A}|}{\delta}\right)}{\varepsilon^2(\alpha)^2\min(\eta^2,1)}\right)$, yields

$$\Pr[\forall \mathbf{a}\in\mathcal{A}, |\hat{c}_{h,\alpha}(\mathbf{a}) - c_\alpha(\mathbf{a})| \le \varepsilon] \ge 1-\delta \ . \tag{33}$$

□

H.9 Proof of proposition 4.4

Proposition 4.4. *If the valuation* $\mathbf{v}$ *is distributed as a multivariate Gaussian, i.e.,* $\mathbf{v}\sim\mathcal{N}(\bar{\mathbf{v}},\Sigma)$*, then, the optimization problem in* (9) *simplifies to*

$$\max_{\mathbf{a}\in\mathcal{A}} \mathbf{a}^{\intercal}\bar{\mathbf{v}} - \frac{\phi(\Phi^{-1}(1-\alpha))}{\alpha}\sqrt{\mathbf{a}^{\intercal}\Sigma\mathbf{a}} \ . \tag{11}$$

Proof. The proof simply follows from the fact that for any normally distributed random variable $X\sim\mathcal{N}(\mu,\sigma^2)$ with mean $\mu\in\mathbb{R}$ and $\sigma\in\mathbb{R}_{0+}$, $CVaR[X] = \mu - \frac{\phi(\Phi(\alpha)^{-1})}{1-\alpha}\sigma$. If the valuation for any group G is normally distributed as $\mathcal{N}(\Sigma_G,\Sigma_G)$, then the utility corresponding to that group has mean $= \frac{\boldsymbol{w}_G}{|G|}\mathbf{a}_G^{\intercal}\bar{\mathbf{v}}_G$ and variance $= \left(\frac{\boldsymbol{w}_G}{|G|}\right)^2\mathbf{a}_G^{\intercal}\Sigma_G\mathbf{a}_G$. Therefore, for the utilitarian welfare objective mean$=\sum_G \frac{\boldsymbol{w}_G}{|G|}\mathbf{a}_G^{\intercal}\bar{\mathbf{v}}_G$ and variance$=\sum_{G\in\mathcal{G}}\left(\left(\frac{\boldsymbol{w}_G}{|G|}\right)^2\mathbf{a}_G^{\intercal}\Sigma_G\mathbf{a}_G\right)$. Substituting by these values in the CVaR formulation for normal random variables, we get the stated results. □

H.10 Linear Program for CVaR of Egalitarian Welfare

Proposition H.3. *Given* h *samples of* $\tilde{\mathbf{v}}$*, i.e.,* $\mathbf{v}^1,\mathbf{v}^2,\mathbf{v}^3,\ldots\mathbf{v}^h$ *sampled from* $\mathcal{D}_{\mathbf{v}}$*, the optimal allocation for the problem in* (12) *can be approximately computed by solving*

$$\begin{aligned}&\max_{\mathbf{a}\in\mathcal{A}}\max_{\boldsymbol{y}\in\mathbb{R}^m_{0+},b\in\mathbb{R}}\left(b-\frac{1}{\alpha}\sum_{j=1}^{M}\boldsymbol{y}_j\right)\\&\forall j\in[1,h],\forall G\in\mathcal{G}:\ \boldsymbol{y}_j\ge\frac{1}{h}\left(b-\frac{1}{|G|}\cdot\mathbf{a}_G^{\intercal}\mathbf{v}_G^j\right)\ .\end{aligned}\tag{34}$$

Proof. Consider the CVaR of egalitarian welfare optimization problem, given by

$$\max_{\mathbf{a}\in\mathcal{A}}\mathrm{CVaR}_\alpha\left[\min_{G\in\mathcal{G}}\frac{1}{|G|}\cdot\mathbf{a}_G^{\intercal}\tilde{\mathbf{v}}_G\right]\quad\equiv\max_{w\in\mathbb{R},\mathbf{a}\in\mathcal{A}}\left\{w-\frac{1}{\alpha}\mathbb{E}\left[\left(w-\min_{G\in\mathcal{G}}\frac{1}{|G|}\cdot\mathbf{a}_G^{\intercal}\tilde{\mathbf{v}}_G\right)_+\right]\right\}\ .\tag{35}$$

Substituting the expectation in the above problem with the empirical expectation computed from the h samples of the valuation matrices, we get

$$\sim\max_{w\in\mathbb{R},\mathbf{a}\in\mathcal{A}}\left\{w-\frac{1}{\alpha}\frac{1}{h}\sum_{i=1}^{M}\left(w-\min_{G\in\mathcal{G}}\frac{1}{|G|}\cdot\mathbf{a}_G^{\intercal}\mathbf{v}_G^i\right)_+\right\}\tag{36}$$

Introducing slack variables $\boldsymbol{y}\in\mathbb{R}^m$, we can write the above problem as

$$\begin{aligned}&\max_{\mathbf{a}\in\mathcal{A}}\max_{\boldsymbol{y}\in\mathbb{R}^h,b\in\mathbb{R}}\left(b-\frac{1}{\alpha}\sum_{j=1}^{m}\boldsymbol{y}_j\right)\\&\forall j\in[1,h]:\ \boldsymbol{y}_j\ge 0\\&\forall j\in[1,h]:\ \boldsymbol{y}_j\ge\frac{1}{h}\left(b-\min_{G\in\mathcal{G}}\frac{1}{|G|}\cdot\mathbf{a}_G^{\intercal}\mathbf{v}_G^j\right)\ .\end{aligned}\tag{37}$$

Without loss of generality, we can represent the above problem as

$$\begin{aligned}
&\max_{\mathbf{a}\in\mathcal{A}} \max_{\boldsymbol{y}\in\mathbb{R}^m, b\in\mathbb{R}} \left(b - \frac{1}{\alpha}\sum_{j=1}^{m} \boldsymbol{y}_j\right) \\
&\forall j \in [1,h] : \boldsymbol{y}_j \geq 0 \\
&\forall j \in [1,h], G \in \mathcal{G} : \boldsymbol{y}_j \geq \frac{1}{h}\left(b - \frac{1}{|G|}\cdot \mathbf{a}_G^\intercal \mathbf{v}_G^j\right) \quad .
\end{aligned} \tag{38}$$

□

H.11 Proof of proposition C.1

Proposition C.1. *If $\mathbf{v}_1, \mathbf{v}_2, \cdots, \mathbf{v}_g$ are $i.i.d$ and normally distributed, i.e., $\forall G \in \mathcal{G}$, $\mathbf{v}_G \sim \mathcal{N}(\bar{\mathbf{v}}_G, \Sigma_G)$, then, the optimization problem in* (12) *simplifies to*

$$\begin{aligned}
&\max_{\mathbf{a}\in\mathcal{A}, t\in\mathbb{R}} t \\
&\textit{s.t. } \forall G \in \mathcal{G} : \left(\frac{1}{|G|}\cdot \mathbf{a}_G^\intercal \bar{\mathbf{v}}_G - t\right)^2 \geq \left(\frac{1}{|G|}\cdot \frac{\phi(\Phi^{-1}(1-\alpha))}{(\alpha)}\right)^2 \mathbf{a}_G^\intercal \Sigma_G \mathbf{a}_G \\
&\forall G \in \mathcal{G} : \left(\frac{1}{|G|}\cdot \mathbf{a}_G^\intercal \bar{\mathbf{v}}_G - t\right) \geq 0 \quad .
\end{aligned} \tag{13}$$

Proof. The proof simply follows from the fact that for any normally distributed random variable $X \sim \mathcal{N}(\mu, \sigma^2)$ with mean $\mu \in \mathbb{R}$ and $\sigma \in \mathbb{R}_{0+}$, $CVaR[X] = \mu - \frac{\phi(\Phi(1-\alpha)^{-1})}{\alpha}\sigma$. If the valuation for any group G is normally distributed as $\mathcal{N}(\Sigma_G, \Sigma_G)$, then the utility corresponding to that group has mean $= \frac{\boldsymbol{w}_G}{|G|}\mathbf{a}_G^\intercal \bar{\mathbf{v}}_G$ and variance $= \left(\frac{\boldsymbol{w}_G}{|G|}\right)^2 \mathbf{a}_G^\intercal \Sigma_G \mathbf{a}_G$. Substituting these values in (34), we get

$$\max_{\mathbf{a}\in\mathcal{A}} \min_{G\in\mathcal{G}} \left(\frac{\boldsymbol{w}_G}{|G|}\cdot \mathbf{a}_G^\intercal \bar{\mathbf{v}}_G\right) - \frac{\boldsymbol{w}_G}{|G|}\cdot \frac{\phi(\Phi^{-1}(\alpha))}{(1-\alpha)} \sqrt{\sum_{G\in\mathcal{G}} \mathbf{a}_G^\intercal \Sigma_G \mathbf{a}_G} \quad . \tag{39}$$

Introducing a slack variable t to represent a lower bound on the group utilities and rearranging the terms we get, we get

$$\begin{aligned}
&\max_{\mathbf{a}\in\mathcal{A}, t\in\mathbb{R}} t \\
&\forall G \in \mathcal{G} : \left(\frac{1}{|G|}\cdot \mathbf{a}_G^\intercal \bar{\mathbf{v}}_G - t\right) \geq \left(\frac{1}{|G|}\cdot \frac{\phi(\Phi^{-1}(\alpha))}{(1-\alpha)}\right) \sqrt{\mathbf{a}_G^\intercal \Sigma_G \mathbf{a}_G} \\
&\forall G \in \mathcal{G} : \left(\frac{1}{|G|}\cdot \mathbf{a}_G^\intercal \bar{\mathbf{v}}_G - t\right) \geq 0 \quad .
\end{aligned} \tag{40}$$

Squaring the quadratic constraint on both sides and adding a constraint to ensure that the non-negativity of the L.H.S of each each group constraint in (40) is retained after squaring, gives us the final result. □

I Details of Code and Machine Specifications

I.1 Code

We have provided the code in the supplementary materials.

I.1.1 Machine Specifications and Computational Time

All experiments were run on Xeon E5-2680 v4 @ 2.40GHz machines with 128GB RAM with each experiment consuming at most 32 GB of memory. We ran 1500 experiments in total and each experiment took 3-4 hours.