Preference Elicitation and Aggregation to Aid with Patient Triage during the COVID-19 Pandemic

Caroline M. Johnston [1] **Simon Blessenohl** [2] **Phebe Vayanos** [1]

Abstract

During the COVID-19 pandemic, triage committees must make ethically difficult decisions which are complicated by the diversity of stakeholder interests. We propose an automated approach to support group decisions by recommending a policy to the group that strikes a compromise between potentially conflicting individual preferences. To identify a policy that best aggregates individual preferences, our system first elicits individual stakeholder value judgements by asking a moderate number of strategically selected queries, each taking the form of a pairwise comparison posed to a specific stakeholder. We propose a novel multi-stage robust optimization formulation of this problem that selects queries that best inform the downstream recommendation problem. Formulating this as a mixed-integer linear program, we evaluate the performance of our approach on the problem of recommending policies for allocating critical care beds to patients with COVID-19. We show that asking questions intelligently allows recommending a policy with significantly lower regret than asking questions randomly, suggesting that the system can help committees reach a better decision by suggesting a policy that aligns with stakeholder value judgments.

1. Introduction

1.1. Motivation

During the COVID-19 pandemic, many hospitals are experiencing a shortage of resources critical to patient care such as ventilators, N-95 masks, or critical care beds (Ranney et al., 2020). Without policies to guide their decisions, doctors face the burden of repeatedly making life-or-death decisions. This not only causes psychological distress in doctors (Greenberg et al., 2020; Ferraresi, 2020), but can also lead to inconsistent and inefficient allocation. Therefore, a preferred and widely adopted approach is to let committees decide on policies for allocating scarce medical resources. Doctors then follow those policies, possibly with the help of on-the-ground triage committees (Emanuel et al., 2020).

But making a collective decision between different policies is no easy task, since it requires making difficult moral trade-offs between, for instance, saving the most lives and giving people equal chances of treatment. Moreover, stakeholders on committees are likely to disagree about which trade-offs are appropriate.

To help committees reach a good decision, we propose a system which learns stakeholder preferences by asking pairwise comparison questions and then recommends an alternative that best aligns with the diverse stakeholder value judgements. The recommended alternative could be directly adopted as the group's decision. More realistically, it could serve as a starting point for further deliberation and decision-making by the group. Finding such an alternative is a challenging problem because asking all possible questions to all individuals to obtain full knowledge about their preferences would impose unreasonable burdens on participants. Thus, the system must recommend policies under uncertainty about stakeholder preferences and intelligently choose which questions to ask to be able to make a good recommendation.

1.2. Background & Literature Review

Much of the literature on preference aggregation and elicitation under uncertainty models agents as having ordinal preference rankings (Conitzer & Sandholm, 2002; Konczak & Lang, 2005; Lu & Boutilier, 2011; Naamani-Dery et al., 2015; Benabbou & Perny, 2016). In contrast, we assume that agents' preferences are determined by cardinal utility functions and that the utilities of different agents lie on the same scale, permitting interpersonal utility comparisons (see (Sen, 1970) for a defense of this assumption). This richer information allows using more nuanced ways of ranking and scoring alternatives for the group, such as summing the agents' utilities. Optimal elicitation and recommen-

[1]Center for Artificial Intelligence in Society, University of Southern California, Los Angeles, CA, USA [2]School of Philosophy, University of Southern California, Los Angeles, CA, USA. Correspondence to: Caroline Johnston <cmjohnst@usc.edu>.

dation based on partial information have been studied in this utility-based setting, but primarily in the single-agent case. Some adopt a Bayesian approach (Chajewska et al., 2000; Boutilier, 2002), while others use a non-probabilistic model in which a set of feasible utility functions is narrowed by discarding utility functions that are inconsistent with the preferences revealed by the agent's responses. In the latter line work, it is often assumed that alternatives are represented by feature vectors and agents' utilities are linear functions in those features or have a generalized additive form (Toubia et al., 2003; 2004; Boutilier et al., 2006; Bertsimas & O'Hair, 2013; Vayanos et al., 2020). Our contribution fits into that line of work, as we also work with sets of feasible utility functions and share the assumption that utility functions are linear. The work closest to our own by (Vayanos et al., 2020), like us, integrates the elicitation and recommendation phases in a single problem to compute optimal queries and recommendations based on partial information. However, we differ from existing contributions in this literature by considering the problem of elicitation and recommendation to a group of agents rather than a single agent. Those who have studied the multi-agent setting with utility-based scoring of alternatives either do not consider optimal elicitation for recommendation (Boutilier et al., 2015; Ferrara et al., 2017) or adopt a probabilistic approach (Zhao et al., 2018).

Our contributions are threefold: a) We propose the first (to the best of our knowledge) formal mathematical formulation of eliciting and aggregating preferences that integrates the learning and recommendation phases under the polyhedral approach to uncertainty modeling; b) We show that this problem, which is sequential and involves uncertainty, can be rewritten equivalently as a mixed-integer linear program (MILP) which can be solved with off-the-shelf solvers; c) We apply our approach to the problem of recommending policies for critical care bed allocation during the COVID-19 pandemic. We show that asking well-chosen questions allows the system to recommend a policy to the group with a much lower sum-utility regret than the best policy one could recommend if one asked questions at random. A small number of queries suffices to significantly reduce the worst-case regret. The system therefore appears to be suited to help a committee reach a better decision by suggesting a promising policy without requiring stakeholders to spend large amounts of time answering queries.

2. Model

We assume there are A agents for which we aim to recommend a single alternative from a set of alternatives $\mathcal{R}$. Before this recommendation, we can ask agents a moderate number K of pairwise comparison queries of the form 'Do you prefer alternative A or alternative B?'.

More formally, let $\mathcal{R} := \{\boldsymbol{x}^i\}_{i\in\mathcal{I}} \subseteq \mathbb{R}^J$, where $\mathcal{I} := \{1,\dots,I\}$ is the set of alternatives and J the number of features of each alternative. A query is a comparison between two alternatives. The set of possible queries is $\mathcal{C} := \left\{(i,i') \in \mathcal{I}^2 \mid i < i'\right\}$. A particular choice of K queries indexed in the set $\mathcal{K} := \{1,\dots,K\}$ is represented by two vectors. First, $\boldsymbol{\iota} \in \mathcal{C}^K$ specifies which alternatives are compared in the K queries. For $\kappa \in \mathcal{K}$, we write $\boldsymbol{\iota}^\kappa := (\boldsymbol{\iota}_1^\kappa, \boldsymbol{\iota}_2^\kappa) \in \mathcal{C}$ to denote that the κth query elicits the preference between alternatives $\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa}$ and $\boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}$. Second, $\boldsymbol{\alpha} \in \mathcal{A}^K$ specifies which agents the different queries are posed to, where $\mathcal{A} := \{1,\dots,A\}$ is the set of agents. Thus, for $\kappa \in \mathcal{K}$, $\boldsymbol{\alpha}_\kappa$ is the agent to whom the κth query is directed.

Each agent's utility function is assumed to be linear in $\boldsymbol{x} \in \mathbb{R}^J$. Thus, it can be represented as a vector of coefficients in $\mathbb{R}^J$. We use $\mathcal{U}^a \subseteq \mathbb{R}^J$ to denote the agent's uncertainty set, the set of feasible utility functions for agent a. Each element $\boldsymbol{u}^a \in \mathcal{U}^a$ represents one possible realization of the agent's utility function. We assume that $\mathcal{U}^a$ is a non-empty bounded polyhedron such that $\mathcal{U}^a = \left\{\boldsymbol{u}^a \in \mathbb{R}^J \mid \boldsymbol{B}^a\boldsymbol{u}^a \geq \boldsymbol{b}^a\right\}$, for some $\boldsymbol{B}^a \in \mathbb{R}^{M\times J}$, $\boldsymbol{b}^a \in \mathbb{R}^M$, a common assumption in the literature (see (Toubia et al., 2003; 2004; Boutilier et al., 2006; Bertsimas & O'Hair, 2013; Vayanos et al., 2020)).

When asked the κth query, an agent is able to respond in one of two ways, using the elements of $\mathcal{S} := \{-1,1\}$: either the agent prefers alternative 1 ($\boldsymbol{s}_\kappa = 1$) or alternative 2 ($\boldsymbol{s}_\kappa = -1$). Our problem formulation assumes that agents are never indifferent. This assumption is innocuous because the resulting problem has the same solutions and objective values as the problem which allows for indifference, as an argument similar to the one given by (Vayanos et al., 2020) shows. After asking agents a series of queries and observing their responses, the updated uncertainty set for agent a, which contains the utility functions consistent with the observed responses, is defined as

$$\mathcal{U}^a(\boldsymbol{\alpha},\boldsymbol{\iota},\boldsymbol{s}) := \left\{\begin{array}{l} \boldsymbol{u}\in\mathcal{U}^a \;:\; \forall\kappa\in\mathcal{K} : \boldsymbol{\alpha}_\kappa = a \\ \boldsymbol{u}^\top\left(\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}\right) \geq 0,\ \boldsymbol{s}_\kappa = 1 \\ \boldsymbol{u}^\top\left(\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}\right) \leq 0,\ \boldsymbol{s}_\kappa = -1 \end{array}\right\}. \tag{1}$$

Given a vector of queries $\boldsymbol{\iota}$ posed to agents $\boldsymbol{\alpha}$, we denote the set of responses consistent with at least one realization of the utility functions in $\mathcal{U}^a$, for each a, by $\mathcal{S}(\boldsymbol{\alpha},\boldsymbol{\iota}) := \left\{\boldsymbol{s}\in\mathcal{S}^K \mid \mathcal{U}^a(\boldsymbol{\alpha},\boldsymbol{\iota},\boldsymbol{s}) \neq \emptyset,\ \forall\, a\in\mathcal{A}\right\}$.

3. Problem Formulation

In "offline elicitation," the decision-maker asks all K queries at once before receiving any of the agents' responses. This scenario arises in situations where preference information is gathered in decentralized systems, such as paper surveys or surveying agents located in different hospitals across the nation.

In "online elicitation," queries are asked one at a time, using previous responses to guide which queries are asked next. This arises in settings such as surveys that are administered by a centralized computer. By incorporating previous responses into the elicitation process, the decision-maker can ask queries providing richer preference information and reducing uncertainty more than in the offline setting.

For both settings, we consider the problem of a decision-maker who aggregates individual utilities for an alternative by summing them (often called "utilitarianism"). To account for the uncertainty in each agent's preferences, the decision-maker recommends an alternative that is robust in the sense that it minimizes the worst-case regret across all utility functions consistent with the agents' responses.

Offline Setting To formulate the offline problem as an optimization problem, note that the decision-maker must first select K queries $\boldsymbol{\iota}^\kappa \in \mathcal{C}$ and agents $\boldsymbol{\alpha}_\kappa \in \mathcal{A}$ for $\kappa \in \mathcal{K}$. The appropriate agents then respond to these queries with answers $\boldsymbol{s}_\kappa$ such that they are consistent with a realization of their utility function in $\mathcal{U}^a$. The decision-maker then recommends an alternative with the minimum worst-case group regret across all possible utility functions of the agents that are consistent with their responses. This problem takes the following form:

$$\min_{\substack{\boldsymbol{\iota}\in\mathcal{C}^K,\\ \boldsymbol{\alpha}\in\mathcal{A}^K}} \max_{\boldsymbol{s}\in\mathcal{S}(\boldsymbol{\alpha},\boldsymbol{\iota})} \min_{\boldsymbol{x}\in\mathcal{R}} \max_{\substack{\boldsymbol{u}^a\in\mathcal{U}^a(\boldsymbol{\alpha},\boldsymbol{\iota},\boldsymbol{s}),\\ \forall a\in\mathcal{A},\\ \boldsymbol{x}'\in\mathcal{R}}} \sum_{a\in\mathcal{A}} \left[(\boldsymbol{u}^a)^\top(\boldsymbol{x}'-\boldsymbol{x})\right]. \tag{2}$$

Online Setting We now informally describe the online procedure. Consider the same scenario as the offline setting, but now the decision-maker selects K queries, $\boldsymbol{\iota}^\kappa \in \mathcal{C}$ and agents $\boldsymbol{\alpha}_\kappa \in \mathcal{A}$ for $\kappa \in \mathcal{K}$ one at a time. Only after observing the answer $\boldsymbol{s}_\kappa$ to the κth query is the $(\kappa+1)$st query decided and the agent's uncertainty set is updated. After receiving answers to all queries that are consistent with each agent's uncertainty set, the decision-maker recommends an alternative that minimizes the regret of the group.

4. MILP Reformulation

In order to model Problem (2) as a finite optimization problem, we allow the recommended alternative to depend on the response scenario $\boldsymbol{s}$ and interchange the inner maximization and minimization terms in Problem (2). Additionally, we can replace $\mathcal{S}(\boldsymbol{\alpha},\boldsymbol{\iota})$ with $\mathcal{S}^K$. This intuitively holds true because, when allowing for inconsistent agent responses, we have that $\mathcal{U}^a(\boldsymbol{\alpha},\boldsymbol{\iota},\boldsymbol{s}) = \emptyset$ for some $a \in \mathcal{A}$ and Problem (2) has an optimal value of $-\infty$. Thus, we could never obtain a solution with a higher objective function value than those of consistent responses. Hence, we can formulate Problem (2) as the following equivalent problem

$$\min_{\substack{\boldsymbol{\iota}\in\mathcal{C}^K,\\ \boldsymbol{\alpha}\in\mathcal{A}^K}} \min_{\substack{\boldsymbol{x}^{\boldsymbol{s}}\in\mathcal{R}:\\ \boldsymbol{s}\in\mathcal{S}^K}} \max_{\boldsymbol{s}\in\mathcal{S}^K} \max_{\substack{\boldsymbol{u}^a\in\mathcal{U}^a(\boldsymbol{\alpha},\boldsymbol{\iota},\boldsymbol{s}),\\ \forall a\in\mathcal{A},\\ \boldsymbol{x}'\in\mathcal{R}}} \sum_{a\in\mathcal{A}} \left[(\boldsymbol{u}^a)^\top(\boldsymbol{x}'-\boldsymbol{x}^{\boldsymbol{s}})\right], \tag{3}$$

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

Now consider the inner maximization of Problem (3) for fixed $\boldsymbol{\iota} \in \mathcal{C}^K$, $\boldsymbol{\alpha} \in \mathcal{A}^K$, $\boldsymbol{x}^{\boldsymbol{s}} \in \mathcal{R} : \boldsymbol{s} \in \mathcal{S}^K$, $\boldsymbol{s} \in \mathcal{S}^K$, $\boldsymbol{x}' \in \mathcal{R}$ using (1) and the polyhedral representation of $\mathcal{U}^a$:

$$\begin{array}{lll} \max & \displaystyle\sum_{a\in\mathcal{A}} \left[(\boldsymbol{u}^a)^\top(\boldsymbol{x}'-\boldsymbol{x}^{\boldsymbol{s}})\right] & \\ \text{s.t.} & (\boldsymbol{u}^a)^\top \boldsymbol{s}_\kappa \left(\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}\right) \geq 0 & \forall \kappa \in \mathcal{K} : \boldsymbol{\alpha}_\kappa = a \\ & \boldsymbol{B}^a \boldsymbol{u}^a \geq \boldsymbol{b}^a & \forall a \in \mathcal{A}. \end{array} \tag{4}$$

By using standard robust optimization techniques, see e.g., (Ben-Tal et al., 2009), we take the dual of Problem (4) and see that Problem (3) is equivalent to the following finite program

$$\begin{array}{ll} \min & \tau \\ \text{s.t.} & \tau \in \mathbb{R}, \boldsymbol{\iota} \in \mathcal{C}^K, \boldsymbol{\alpha} \in \mathcal{A}^K, \boldsymbol{x}^{\boldsymbol{s}} \in \mathcal{R}, \forall \boldsymbol{s} \in \mathcal{S}^K \\ & \left.\begin{array}{l} \boldsymbol{\zeta}^{(\boldsymbol{x}',\boldsymbol{s})} \in \mathbb{R}_-^K \\ \boldsymbol{\beta}^{(\boldsymbol{x}',\boldsymbol{s},a)} \in \mathbb{R}_-^M \\ \displaystyle\sum_{\substack{\kappa\in\mathcal{K}:\\ \boldsymbol{\alpha}_\kappa = a}} \boldsymbol{s}_\kappa \left(\boldsymbol{x}^{\boldsymbol{\iota}_1^\kappa} - \boldsymbol{x}^{\boldsymbol{\iota}_2^\kappa}\right) \zeta_\kappa^{(\boldsymbol{x}',\boldsymbol{s})} \\ \quad +(\boldsymbol{B}^a)^\top \boldsymbol{\beta}^{(\boldsymbol{x}',\boldsymbol{s},a)} = \boldsymbol{x}' - \boldsymbol{x}^{\boldsymbol{s}} \end{array}\right\} \begin{array}{l} \forall \boldsymbol{x}' \in \mathcal{R}, \\ \forall \boldsymbol{s} \in \mathcal{S}^K, \\ \forall a \in \mathcal{A} \end{array} \\ & \tau \geq \displaystyle\sum_{a\in\mathcal{A}} (\boldsymbol{\beta}^{(\boldsymbol{x}',\boldsymbol{s},a)})^\top \boldsymbol{b}^a \quad \forall \boldsymbol{x}' \in \mathcal{R}, \forall \boldsymbol{s} \in \mathcal{S}^K, \end{array} \tag{5}$$

where $\boldsymbol{\zeta}^{(\boldsymbol{x}',\boldsymbol{s})}$ and $\boldsymbol{\beta}^{(\boldsymbol{x}',\boldsymbol{s},a)}$ are the set of dual variables corresponding to the first (resp. second) set of constraints in Problem (4). Problem (5) can then be converted to an MILP by introducing binary variables to encode the choice of $\boldsymbol{\alpha}$ and $\boldsymbol{\iota}$ and using standard linearization techniques, see e.g., (Hillier & Lieberman, 2001).

This MILP is solved directly for the offline problem. For the online problem, we solve a series of MILPs, using a conservative approximation with a folding horizon approach. For each period, we ask the query that is optimal if this were the final query to be asked and update the uncertainty set accordingly. At the end of the planning horizon, we make the recommendation that is robust against any utility vector that is consistent with the responses collected over time.

5. Resource Allocation during the Pandemic

We apply our approach to the problem of scarce resource allocation during the COVID-19 pandemic (see Section 1.1).

5.1. Simulating Policy Outcomes

We simulate how different policies assign scarce critical care beds to COVID-19 patients. The metrics of a policy's performance, or the features by which agents evaluate them, are the total number of life-years saved, the probabilities of receiving critical care across different age groups, and the survival probabilities across different age groups.

For each day, we simulate the arrival of patients, the assignment of waiting patients to free critical care beds by the policy based on patient characteristics, as well as the recovery or death of patients. We run this simulation at the country-level, with UK data, for April 1st to July 15th, 2020. To set simulation parameters, we use projections of the Institute for Health Metrics and Evaluation (IHME) model[1] and outcome data from patients with COVID-19 from the UK Intensive Care National Audit and Research Centre.[2] We only use age as a patient characteristic. The primary obstacle to using more characteristics is the unavailability of patient outcome data based on, say, both age and race.

A policy assigns a score to each patient based on their age and the number of days they waited for a critical care bed. It then allocates free beds to highest-scoring patients. We consider policies that use regression trees to assign a score to a patient. Each node contains a condition on the patient's age or waiting time, and all patients that reach the same leaf are assigned the same score. We generate 25 regression trees of depth 3 by randomly picking a feature and a comparison value for each non-leaf node, and a random number between 0 and 1 for leaf nodes. We then simulate the outcomes of each policy, obtaining 25 alternatives with 15 features each. Therefore, there are $25 \cdot 24 = 600$ possible pairwise comparison questions each agent could be asked.

5.2. Preference Elicitation Results

For both elicitation types we assume the following scenario. Using the policies in Section 5.1, we solve the MILP formulation of Problem (5) (or a series of the MILP in the online case) with $\mathcal{A} = \{1, 2\}$. We assume that each component of the utility vector for both agents lies in $[0, 1]$, i.e., they have a non-negative utility for each policy feature. Finally, we normalize the worst case regret between the scenario in which no questions are asked (regret value of 1) and a conservative lower bound in which we have complete knowledge of the agents' utility functions (regret value of 0).

Offline Elicitation Using a decomposition technique and a conservative approximation approach to speed up computations for solving the MILP (see (Vayanos et al., 2020)), we show results for the worst case regret over simulated policies in Figure 1 for $K = 10$. Our method (`MinMaxRegret`) outperforms asking random queries (`Random`).

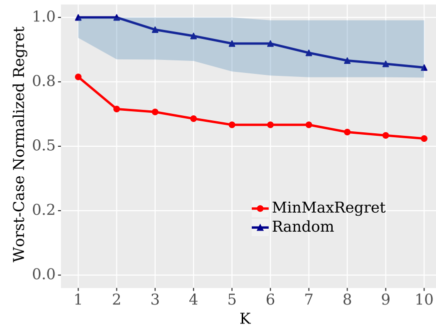

Figure 1. Optimality results for offline elicitation for $\mathcal{A} = \{1, 2\}$. The blue shaded region corresponds to asking random queries over 10 instances, where the blue line corresponds to the median performance. The red line corresponds to the performance of asking $\min \max$ regret queries over a single instance.

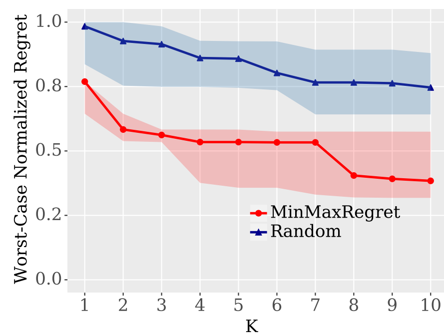

Figure 2. Optimality results for online elicitation for $\mathcal{A} = \{1, 2\}$. The blue (resp. red) shaded region corresponds to asking random (resp. $\min \max$ regret) queries over 10 instances of 2 agents. The red and blue line correspond to the median performance of each.

Online Elicitation To simulate the responses of agents, we sample their true utility vectors uniformly at random from the non-negative values in a J-dimensional sphere of radius one. The optimality results over 10 instances of 2 agents are shown in Figure 2 for $K = 10$. We see that again `MinMaxRegret` outperforms `Random`; our method's worst performing instance obtains a lower regret value than the best performing random instance. On average, online elicitation leads to lower regret than offline elicitation, obtaining more utility knowledge with fewer queries.

6. Future Work

Using our preference elicitation framework, stakeholders can streamline the collective decision-making process by providing a compromise solution in settings such as scarce

[1] http://www.healthdata.org/covid/data-downloads
[2] https://www.icnarc.org/Our-Audit/Audits/Cmp/Reports

resource allocation during the COVID-19 pandemic. In future work, we aim to explore more realistic committee sizes, different methods of aggregating individual preferences, and the extent to which the suggested approach is incentive-compatible, in the sense that agents do not have an incentive to misreport their preferences.

7. Acknowledgements

We are grateful to the authors of (Vayanos et al., 2020), whose code (available at https://github.com/duncanmcelfresh/ActiveRobustPreferenceElicitation) we used as a starting point for our simulations in Section 5.

References

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

Benabbou, N. and Perny, P. Solving multi-agent knapsack problems using incremental approval voting. In *Proceedings of the Twenty-Second European Conference on Artificial Intelligence*, ECAI'16, pp. 1318–1326, 2016.

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

Boutilier, C. A pomdp formulation of preference elicitation problems. In *Proceedings of The Eighteenth National Conference On Artificial Intelligence*, pp. 239–246, 2002.

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

Boutilier, C., Caragiannis, I., Haber, S., Lu, T., Procaccia, A. D., and Sheffet, O. Optimal social choice functions: A utilitarian view. *Artificial Intelligence*, 227:190–213, 2015.

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

Conitzer, V. and Sandholm, T. Vote elicitation: Complexity and strategy-proofness. In *AAAI/IAAI*, pp. 392–397, 2002.

Emanuel, E. J., Persad, G., Upshur, R., Thome, B., Parker, M., Glickman, A., Zhang, C., Boyle, C., Smith, M., and Phillips, J. P. Fair allocation of scarce medical resources in the time of covid-19. *New England Journal of Medicine*, 382(21):2049–2055, 2020.

Ferrara, M., Rasouli, S., Khademi, M., and Salimi, M. A robust optimization model for a decision-making problem: An application for stock market. *Operations Research Perspectives*, 4:136–141, 2017. ISSN 22147160. doi: 10.1016/j.orp.2017.10.001. URL https://doi.org/10.1016/j.orp.2017.10.001.

Ferraresi, M. A coronavirus cautionary tale from italy: Don't do what we did. *The Boston Globe*, 2020. URL https://www.bostonglobe.com/2020/03/13/opinion/coronavirus-cautionary-tale-italy-dont-do-what-we-did/.

Greenberg, N., Docherty, M., Gnanapragasam, S., and Wessely, S. Managing mental health challenges faced by healthcare workers during covid-19 pandemic. *BMJ*, 368, 2020.

Hillier, F. S. and Lieberman, G. J. *Introduction to Operations Research*. McGraw-Hill, New York, NY, USA, seventh edition, 2001.

Konczak, K. and Lang, J. Voting procedures with incomplete preferences. In *Proceedings of the IJCAI-05 Multidisciplinary Workshop on Advances in Preference Handling*, volume 20, 2005.

Lu, T. and Boutilier, C. Robust approximation and incremental elicitation in voting protocols. In *Twenty-Second International Joint Conference on Artificial Intelligence*, 2011.

Naamani-Dery, L., Golan, I., Kalech, M., and Rokach, L. Preference elicitation for group decisions using the borda voting rule. *Group Decision and Negotiation*, 24(6):1015–1033, 2015.

Ranney, M. L., Griffeth, V., and Jha, A. K. Critical supply shortages—the need for ventilators and personal protective equipment during the covid-19 pandemic. *New England Journal of Medicine*, 2020.

Sen, A. *Collective Choice and Social Welfare*. Holden-Day, San Francisco, 1970.

Toubia, O., Simester, D. I., Hauser, J. R., and Dahan, E. Fast polyhedral adaptive conjoint estimation. *Marketing Science*, 22(3):273–303, 2003.

Toubia, O., Hauser, J. R., and Simester, D. I. Polyhedral methods for adaptive choice-based conjoint analysis. *Journal of Marketing Research*, 41(1):116–131, 2004.

Vayanos, P., McElfresh, D. C., Ye, Y., Dickerson, J. P., and Rice, E. Active preference elicitation via adjustable robust optimization. 2020. Available at https://arxiv.org/abs/2003.01899.

Zhao, Z., Li, H., Wang, J., Kephart, J., Mattei, N., Su, H., and Xia, L. A cost-effective framework for preference elicitation and aggregation. *arXiv preprint arXiv:1805.05287*, 2018.