

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

Caroline M. Johnston,¹ Simon Blessenohl,² Phebe Vayanos¹

¹University of Southern California, Center for Artificial Intelligence in Society

²University of Southern California, School of Philosophy

{cmjohnst, blesseno, phebe.vayanos}@usc.edu

Abstract

During the COVID-19 pandemic, committees have been appointed to make ethically difficult triage decisions, which are complicated by the diversity of stakeholder interests involved. We propose a disciplined, automated approach to support such difficult collective decision-making. Our system aims to recommend 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 formulation of this problem that selects which queries to ask which individuals to best inform the downstream recommendation problem. Modeling this as a multi-stage robust optimization problem, we show that we can equivalently reformulate this as a mixed-integer linear program which can be solved with off-the-shelf solvers. 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 the system to recommend a policy with a much lower regret than asking questions randomly. The lower regret suggests that the system is suited to help a committee reach a better decision by suggesting a policy that aligns with stakeholder value judgements.

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 having to repeatedly make life-or-death decisions. This not only causes psychological distress in doctors [Greenberg *et al.*, 2020; Ferraresi, 2020], but could 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 can 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, the stakeholders on the committees are likely to disagree about what trade-offs are appropriate.

To help committees reach a good decision, we propose a system which learns about 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.

While we focus on collective decisions about policies for allocating scarce medical resources, our approach is applicable to any collective decision. For instance, it could also be used to help committees decide between policies for allocating scarce housing resources to those experiencing homelessness.

1.2 Background & Literature Review

There is a rich literature on different rules for mapping information about individual preferences or utilities to a collective ranking or scoring of the alternatives. The most commonly studied problem assumes that the full preference rankings for all agents are known, and investigates preference aggregation functions, which map individual preference rankings to a ranking for the group [Arrow, 1951]. The case of incomplete information about agents’ preference has also been studied [Conitzer and Sandholm, 2002; Konczak and Lang, 2005; Benabbou and Perny, 2016]. This invites investigating which individual preferences to elicit to make a highly-ranked rec-

ommendation based on partial information [Pini *et al.*, 2011; Lu and Boutilier, 2011; Naamani-Dery *et al.*, 2015].

We consider elicitation and recommendation in a slightly different setting. We assume that the agents’ preferences are determined by cardinal utility functions and, furthermore, 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 about agents allows using more nuanced ways of ranking and scoring alternatives for the group, such as summing individual utilities. Questions of optimal elicitation and recommendation based on partial information have been addressed in this utility-based setting, but primarily in the single-agent case. Some adopt a Bayesian approach, in which a prior probability distribution over the agent’s possible utility functions is updated based on the agent’s responses to pairwise comparison questions [Chajewska *et al.*, 2000; Boutilier, 2002]. Others adopt a non-probabilistic approach 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 this 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; Toubia *et al.*, 2004; Boutilier *et al.*, 2006; Bertsimas and 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 in that we consider the problem of elicitation and recommendation based on partial information to a group of agents rather than a single agent. The multi-agent setting has received less attention so far. [Boutilier *et al.*, 2015; Ferrara *et al.*, 2017] consider aggregation with utility-based scoring of alternatives based on ordinal individual preference information, but they do not study how to optimally elicit subsets of preferences with the goal of ultimately recommending an item. [Zhao *et al.*, 2018] consider elicitation and aggregation in a multi-agent setting, but they use a probabilistic approach and voting rules to aggregate individual utilities.

Our contributions are threefold: a) We propose the first (to the best of our knowledge) formal mathematical formulation of the problem of eliciting and aggregating preferences in a way 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 simulate policy outcomes using real data from the UK. 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?’ to guide this decision.

More formally, let $\mathcal{R} := \{\mathbf{x}^i\}_{i \in \mathcal{I}} \subseteq \mathbb{R}^J$, where $\mathcal{I} := \{1, \dots, I\}$ is the set of alternatives. A query is a comparison between two alternatives. The set of possible queries is $\mathcal{C} := \{(i, i') \in \mathcal{I}^2 \mid i < i'\}$. 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 κ th query elicits the preference between alternatives $\mathbf{x}^{\boldsymbol{\iota}_1^\kappa}$ and $\mathbf{x}^{\boldsymbol{\iota}_2^\kappa}$. 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 κ th query is directed.

Each agent’s utility function is assumed to be linear in $\mathbf{x} \in \mathbb{R}^J$. 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 $\mathbf{u}^a \in \mathcal{U}^a$ represents one possible realization of the agent’s utility function. We assume that \mathcal{U}^a is a non-empty bounded polyhedron such that $\mathcal{U}^a = \{\mathbf{u}^a \in \mathbb{R}^J \mid \mathbf{B}^a \mathbf{u}^a \geq \mathbf{b}^a\}$, for some $\mathbf{B}^a \in \mathbb{R}^{M \times J}$, $\mathbf{b}^a \in \mathbb{R}^M$, a very common assumption in the literature (see [Toubia *et al.*, 2003; Toubia *et al.*, 2004; Boutilier *et al.*, 2006; Bertsimas and O’Hair, 2013; Vayanos *et al.*, 2020]).

When asked the κ th 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 ($s_\kappa = 1$) or prefers alternative 2 ($s_\kappa = -1$). 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

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

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 $\mathcal{S}(\boldsymbol{\alpha}, \boldsymbol{\iota}) := \{s \in \mathcal{S}^K \mid \mathcal{U}^a(\boldsymbol{\alpha}, \boldsymbol{\iota}, s) \neq \emptyset, \forall a \in \mathcal{A}\}$.

3 Problem Formulation

In “offline elicitation,” the decision-maker asks all K pairwise comparison 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 locations/time zones such as hospitals across the nation.

In “online elicitation,” queries are asked one at a time, using the responses to previous queries 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 In the offline setting, the decision-maker can select K total pairwise comparison queries to ask of any of these agents, which must be decided before receiving any answers. To formulate this as an optimization problem, note that the decision-maker must first select K queries $\iota^\kappa \in \mathcal{C}$ and agents $\alpha_\kappa \in \mathcal{A}$ for $\kappa \in \mathcal{K}$. The appropriate agents then respond to these queries with answers s_κ such that they are consistent with a realization of their utility function in \mathcal{U}^a . The decision-maker can then confirm that $u^a \in \mathcal{U}^a(\alpha, \iota, s)$ for each agent and then recommends an alternative with the minimum worst-case group regret. This problem takes the following form:

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

Online Setting We now informally describe the online procedure. Consider the same regret scenario as the offline setting, but now the decision-maker selects K queries, $\iota^\kappa \in \mathcal{C}$ and agents $\alpha_\kappa \in \mathcal{A}$ for $\kappa \in \mathcal{K}$ one at a time. Only after observing the answer s_κ to the κ th query is the $(\kappa + 1)$ st query decided and the agent’s uncertainty set is updated. After receiving answers to all K 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 s , allowing us to interchange the inner maximization and minimization terms in Problem (2). Additionally, we can replace $\mathcal{S}(\alpha, \iota)$ with \mathcal{S}^K . This intuitively holds true because, when allowing for inconsistent agent responses, we have that $\mathcal{U}^a(\alpha, \iota, s) = \emptyset$ 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

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

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

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

$$\begin{aligned} \max \quad & \sum_{a \in \mathcal{A}} [(u^a)^\top (x' - x^s)] \\ \text{s.t.} \quad & (u^a)^\top s_\kappa (x^{\iota_1^\kappa} - x^{\iota_2^\kappa}) \geq 0 \quad \forall \kappa \in \mathcal{K} : \alpha_\kappa = a \\ & B^a u^a \geq b^a \quad \forall a \in \mathcal{A}. \end{aligned} \quad (4)$$

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

$$\begin{aligned} \min \quad & \tau \\ \text{s.t.} \quad & \tau \in \mathbb{R}, \iota \in \mathcal{C}^K, \alpha \in \mathcal{A}^K, x^s \in \mathcal{R}, \forall s \in \mathcal{S}^K \\ & \left. \begin{aligned} & \zeta(x', s) \in \mathbb{R}_-^K \\ & \beta(x', s, a) \in \mathbb{R}_-^M \\ & \sum_{\substack{\kappa \in \mathcal{K}: \\ \alpha_\kappa = a}} s_\kappa (x^{\iota_1^\kappa} - x^{\iota_2^\kappa}) \zeta_\kappa(x', s) \\ & + (B^a)^\top \beta(x', s, a) = x' - x^s \end{aligned} \right\} \begin{aligned} & \forall x' \in \mathcal{R}, \\ & \forall s \in \mathcal{S}^K, \\ & \forall a \in \mathcal{A} \end{aligned} \\ & \tau \geq \sum_{a \in \mathcal{A}} (\beta(x', s, a))^\top b^a \quad \forall x' \in \mathcal{R}, \forall s \in \mathcal{S}^K, \end{aligned} \quad (5)$$

where $\zeta(x', s)$ and $\beta(x', s, a)$ 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 α and ι and using standard linearization techniques, see e.g., [Hillier and Lieberman, 2001].

This MILP is solved directly for the offline problem. Using a conservative approximation that uses a folding horizon approach, we solve a series of MILPs for the online problem. For each period, we ask the query that is optimal if this were the final query to be asked and then 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 now 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 will 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) model¹ as well as 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 in our model. The primary obstacle to using more patient 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 the offline and online elicitation we assume the following scenario. Using the policies described in Section 5.1, we solve the MILP formulation of Problem (5) (or a series of the MILP formulation 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, corresponding to a regret of 1, and a conservative lower bound in which we have complete knowledge of the agents’ utility functions, corresponding to a regret of 0.

Offline Elicitation Using both a decomposition technique and a conservative approximation approach to speed up computations while solving the MILP (see [Vayanos *et al.*, 2020]), we show the results for the worst case regret over the simulated policies in Figure 1 for $K = 1, \dots, 10$. We see that our method (MinMaxRegret) outperforms that of asking random queries (Random).

Online Elicitation To simulate the responses of each agent, we sample their true utility vector uniformly at random 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. In fact, our method’s worst performing realization obtains a lower regret value than that of the best performing 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 respective method.

realization, a promising result. Additionally, on average, the online elicitation method leads to lower regret than the offline elicitation for both MinMaxRegret and Random, obtaining more utility information with fewer queries asked.

6 Future Work

Using our preference elicitation framework, stakeholders are able to streamline their collective decision-making process by providing a compromise solution in both socially and time sensitive settings such as scarce resource allocation during the COVID-19 pandemic. For future work and exploration, we aim to increase the number of agents to a more realistic committee group size (e.g., 10 agents) and to explore different methods of aggregating individual preferences, investigating, among other things, the extent to which they provide incentives for agents 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.

¹<http://www.healthdata.org/covid/data-downloads>

²<https://www.icnarc.org/Our-Audit/Audits/Cmp/Reports>

References

- [Arrow, 1951] Kenneth Arrow. *Social Choice and Individual Values*. John Wiley & Sons, New York, 1951.
- [Ben-Tal *et al.*, 2009] A. Ben-Tal, L. El Ghaoui, and A.S. Nemirovski. *Robust Optimization*. Princeton Series in Applied Mathematics. Princeton University Press, October 2009.
- [Benabbou and Perny, 2016] Nawal Benabbou and Patrice Perny. Solving multi-agent knapsack problems using incremental approval voting. In *Proceedings of the Twenty-Second European Conference on Artificial Intelligence, ECAI'16*, page 1318–1326, 2016.
- [Bertsimas and O’Hair, 2013] Dimitris Bertsimas and Allison O’Hair. Learning preferences under noise and loss aversion: An optimization approach. *Operations Research*, 61(5):1190–1199, 2013.
- [Boutilier *et al.*, 2006] 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.
- [Boutilier *et al.*, 2015] Craig Boutilier, Ioannis Caragiannis, Simi Haber, Tyler Lu, Ariel D Procaccia, and Or Sheffet. Optimal social choice functions: A utilitarian view. *Artificial Intelligence*, 227:190–213, 2015.
- [Boutilier, 2002] Craig Boutilier. A pomdp formulation of preference elicitation problems. In *Proceedings of The Eighteenth National Conference On Artificial Intelligence*, pages 239–246, 2002.
- [Chajewska *et al.*, 2000] Urszula Chajewska, Daphne Koller, and Ronald Parr. Making rational decisions using adaptive utility elicitation. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence*, pages 363–369, 2000.
- [Conitzer and Sandholm, 2002] Vincent Conitzer and Thomas Sandholm. Vote elicitation: Complexity and strategy-proofness. In *AAAI/IAAI*, pages 392–397, 2002.
- [Emanuel *et al.*, 2020] Ezekiel J. Emanuel, Govind Persad, Ross Upshur, Beatriz Thome, Michael Parker, Aaron Glickman, Cathy Zhang, Connor Boyle, Maxwell Smith, and James P. Phillips. Fair allocation of scarce medical resources in the time of covid-19. *New England Journal of Medicine*, 382(21):2049–2055, 2020.
- [Ferrara *et al.*, 2017] Massimiliano Ferrara, Sepideh Rasouli, Mehrnoosh Khademi, and Mehdi Salimi. A robust optimization model for a decision-making problem: An application for stock market. *Operations Research Perspectives*, 4:136–141, 2017.
- [Ferraresi, 2020] Mattia Ferraresi. A coronavirus cautionary tale from italy: Don’t do what we did. *The Boston Globe*, 2020.
- [Greenberg *et al.*, 2020] Neil Greenberg, Mary Docherty, Sam Gnanapragasam, and Simon Wessely. Managing mental health challenges faced by healthcare workers during covid-19 pandemic. *BMJ*, 368, 2020.
- [Hillier and Lieberman, 2001] Frederick S. Hillier and Gerald J. Lieberman. *Introduction to Operations Research*. McGraw-Hill, New York, NY, USA, seventh edition, 2001.
- [Konczak and Lang, 2005] Kathrin Konczak and Jérôme Lang. Voting procedures with incomplete preferences. In *Proceedings of the IJCAI-05 Multidisciplinary Workshop on Advances in Preference Handling*, volume 20, 2005.
- [Lu and Boutilier, 2011] Tyler Lu and Craig Boutilier. Robust approximation and incremental elicitation in voting protocols. In *Twenty-Second International Joint Conference on Artificial Intelligence*, 2011.
- [Naamani-Dery *et al.*, 2015] Lihi Naamani-Dery, Inon Golan, Meir Kalech, and Lior Rokach. Preference elicitation for group decisions using the borda voting rule. *Group Decision and Negotiation*, 24(6):1015–1033, 2015.
- [Pini *et al.*, 2011] Maria Silvia Pini, Francesca Rossi, Kristen Brent Venable, and Toby Walsh. Incompleteness and incomparability in preference aggregation: Complexity results. *Artificial Intelligence*, 175(7-8):1272–1289, 2011.
- [Ranney *et al.*, 2020] Megan L Ranney, Valerie Griffeth, and Ashish K Jha. Critical supply shortages—the need for ventilators and personal protective equipment during the covid-19 pandemic. *New England Journal of Medicine*, 2020.
- [Sen, 1970] Amartya Sen. *Collective Choice and Social Welfare*. Holden-Day, San Francisco, 1970.
- [Toubia *et al.*, 2003] Olivier Toubia, Duncan I Simester, John R Hauser, and Ely Dahan. Fast polyhedral adaptive conjoint estimation. *Marketing Science*, 22(3):273–303, 2003.
- [Toubia *et al.*, 2004] Olivier Toubia, John R Hauser, and Duncan I Simester. Polyhedral methods for adaptive choice-based conjoint analysis. *Journal of Marketing Research*, 41(1):116–131, 2004.
- [Vayanos *et al.*, 2020] Phebe Vayanos, Duncan C. McElfresh, Yingxiao Ye, John P. Dickerson, and Eric Rice. Active preference elicitation via adjustable robust optimization. 2020. Available at <https://arxiv.org/abs/2003.01899>.
- [Zhao *et al.*, 2018] Zhibing Zhao, Haoming Li, Junming Wang, Jeffrey Kephart, Nicholas Mattei, Hui Su, and Lirong Xia. A cost-effective framework for preference elicitation and aggregation. *arXiv preprint arXiv:1805.05287*, 2018.