

Majorized Bayesian Persuasion and Fair Selection

Siddhartha Banerjee[†] Kamesh Munagala[‡] Yiheng Shen[‡] Kangning Wang[§]

Abstract

We address the fundamental problem of selection under uncertainty by modeling it from the perspective of Bayesian persuasion. In our model, a decision maker with imperfect information always selects the option with the highest expected value. We seek to achieve fairness among the options by revealing additional information to the decision maker and hence influencing its subsequent selection. To measure fairness, we adopt the notion of majorization, aiming at simultaneously approximately maximizing all symmetric, monotone, concave functions over the utilities of the options. As our main result, we design a novel information revelation policy that achieves a logarithmic-approximation to majorization in polynomial time. On the other hand, no policy, regardless of its running time, can achieve a constant-approximation to majorization. Our work is the first non-trivial majorization result in the Bayesian persuasion literature with multi-dimensional information sets.

1	Introduction	1
1.1	Related Work	3
2	Our Model and Results	5
2.1	Signaling Model	5
2.2	Fairness via Approximate Majorization	5
2.3	Our Results	7
3	Warmup: The FULL REVELATION Signaling Policy	9
3.1	1-Majorization for Bernoulli Distributions: Proof of Theorem 2.5	9
3.2	Majorization among FULL INFORMATION Schemes: Proof of Theorem 2.6	12
4	Approximate Majorization: Proof of Theorem 2.8	14
4.1	SINGLE MEAN Projections	14
4.2	Structure of SINGLE MEAN Projections	15
4.3	Reduction to Network Flow and Approximate Majorization	17
4.4	Extension to Asymmetric Utilities	19
5	Lower Bound on Majorization: Proof of Theorem 2.9	20
6	Open Directions	22

[†]School of Operations Research and Information Engineering, Cornell University, Ithaca, NY 14853. Email: sbanerjee@cornell.edu. Supported by NSF grants ECCS-1847393, CNS-195599 and AFOSR grant FA9550-23-1-0068.

[‡]Department of Computer Science, Duke University, Durham, NC 27708-0129. Emails: kamesh@cs.duke.edu, yiheng.shen@duke.edu. Supported by NSF grants CCF-2113798 and IIS-2402823.

[§]Department of Computer Science, Rutgers University, Piscataway, NJ 08854. Email: kn.w@rutgers.edu.

1 Introduction

Selection problems arise in hiring decisions, college admissions, ad auctions, and many other scenarios. Typically, a selection process involves defining an *explainable* selection rule that solves an optimization problem over the agents (e.g., applicants or advertisers). Perhaps the simplest and most common rule is to choose the best or most qualified agent(s) – for example, many countries use standardized test scores for college admissions. Under natural simplifying assumptions, such a meritocratic rule will maximize the utilitarian social welfare, and can be perceived as fair. Nevertheless, its outcomes may not always align with other fairness considerations. For instance, one can adopt the view of max-min fairness and aim at helping the worst-off agent.

The meritocratic rule has a perhaps more subtle issue: when trying to evaluate the agents and define the “best” one, the quality measurements may be imperfect. Indeed, the decision maker uses observable information or features of the agents to reach a decision, but counterfactually, a different decision could have been made in the presence of more certain or complete information. This issue can have disproportionate negative impacts on certain demographic groups and therefore undermine fairness perceptions towards the rule. To address the issue, there has been a long line of recent research focusing on modeling this uncertainty and constructing randomized selection policies to achieve fair outcomes [KR18, CMV20, SKJ21, SWZ⁺23, DKKS24]. As a practical example, in many hiring situations, the interviewers are required to interview at least a certain number of candidates from prespecified demographic groups, as an attempt to mitigate the impact of such uncertainty and to achieve fairness.

Fairness via Information Revelation. We diverge from this line of research and take a view inspired by the influential Bayesian persuasion literature, starting with the seminal work of Kamenica and Gentzkow [KG11]. We ask:

Given uncertainty about agent qualities, how much additional information should be gleaned and revealed to the decision maker (either by the agents themselves or by a central entity), so that even if the decision maker sticks to “selecting the best”, it still yields a fair outcome to the agents?

To approach this high-level question, we use the following model. (See Au and Kawai [AK20] for a related model.) Each agent has an uncertain scalar quality, and its true quality is known to an entity – either the agent itself or an intermediary. This entity then signals additional information independently for each agent to the decision maker to refine its uncertainty via Bayes’ rule. Analogous to [AK20], this signaling process can be done in a distributed fashion by the agents themselves, a property that is often desirable for the applications we consider in order to preserve privacy of agents’ true quality. Subsequently, it can happen that several agents have comparable posterior quality in the eyes of the decision maker, and a fair selection can be made by the decision maker without significantly sacrificing its own optimality notion of selecting the posterior best.

As a concrete motivation, consider designing a standardized form for admission or hiring. The form is the signaling scheme - it is designed centrally by the decision maker, filled separately by agents, and the decision maker can break ties based on information revealed in the forms, sometimes in conjunction with a lottery number assigned to each agent, to ensure fairness in selection. We present the formal model in [Section 2](#), and an example in [Example 2.4](#).

In order to define fairness, we consider the vector of expected utilities received by the agents, where each agent receives utility equal to its true value if it is selected. We seek signaling (or information revelation) policies whose utility vector is *approximately majorized* [HLP34, GMP05, KK06, CS19], meaning

that all symmetric, monotone, concave functions over the utilities are simultaneously approximately maximized by the same policy. Such fairness functions capture, for instance, max-min fairness (maximizing the minimum utility) and the Nash welfare (maximizing the geometric mean of the utilities).

At one extreme, if the agents do not reveal any additional information, the agent with the highest expected quality, according to the decision maker’s prior, is selected. As mentioned before, this can be problematic in terms of fairness to counterfactual information. We show in [Example 2.7](#) that the other extreme – where each agent reveals its exact quality (FULL REVELATION) – may be far from fair as well, even when the qualities only take “high”, “medium”, and “low” values. This example highlights the need to carefully choose which information to reveal.

We note that such information revelation considerations extend beyond the motivating settings described above. As an example, consider a government agency (the decision maker) that has fixed rules for allocating money designated to a social welfare cause such as refugee resettlement. Local agencies typically are better informed about welfare needs of the groups they serve, and can selectively reveal information in order to facilitate fairness for these groups.

Summary of Results. We present a detailed summary of our results in [Section 2](#) after we present our formal model. Informally, our main result is a novel set of signaling (or information revelation) mechanisms that achieve approximate majorization. (See [Theorem 2.8](#).) The key technical hurdle with designing such policies is the behavior of the decision maker – this decision maker chooses the best (or approximately best) agent given its information. For any given fairness objective (such as max-min fairness), a general algorithmic result due to Dughmi and Xu [[DX16](#)] yields an FPTAS for any Bayesian persuasion problem with arbitrary information sets, assuming that the sender can correlate agent signals. However, such a generic approach does not shed light on the existence of policies that *simultaneously* approximate any fairness objective, and this fact requires us to unearth specific structure in our problem.

Our main contribution in [Section 4](#) is in unearthing such structure for the selection problem via a novel class of signaling policies with simple structure, leading to positive *existence* and *computational* results for approximately majorized policies. We term the novel set of policies as **SINGLE MEAN**, and these generate, in a randomized fashion, information signals where the posterior inferred by the decision maker maps the agents’ values to a common mean quality. This mapping to the common mean happens with as large a probability as possible, hence giving the decision maker wide leeway in implementing a fair selection even when it “goes with the best”. Though this policy sounds intuitive, its analysis is far from obvious as we discuss below. Informally, the final result is a polynomial-time-computable and $O(\log V)$ -approximately majorized policy¹, where V is the ratio of the largest to smallest quality value. This means for any fairness function, the policy yields an $O(\log V)$ -approximation in polynomial time.

As a corollary, our work also presents an approximation algorithm for any given fairness function when agents generate signals independently of other agents. As shown in [Section 4.4](#), our results easily extend to handle the case where the agents’ utility (on which we seek to be fair) is different from their quality (whose posterior is used by the social planner to perform allocation). We note that the FPTAS for general Bayesian persuasion in [[DX16](#)] assumes a more powerful intermediary that can send a common signal using the information of all agents.² Our work shows that for the weaker and practically motivated intermediary that generates per-agent signals independently of other agents, there is still a polynomial-time-computable $O(\log V)$ -approximation.

In [Section 5](#), we complement this positive result by showing that a mild dependence on V is unavoidable – no signaling scheme can be $(\log \log V)/3$ -approximately majorized.

¹Our results are bicriteria, and assume the societal decision maker is a $(1 + \varepsilon)$ -approximate welfare maximizer.

²They also present an independent signaling scheme when the actions have i.i.d. rewards. In our case, the i.i.d. setting is uninteresting since revealing full information about agent quality is trivially 1-majorized.

Technical Highlight: Reduction to Majorized Network Flow. At a technical level, our majorization results proceed by reducing the problem to network flow with a source and multiple sinks, for which a majorized solution has long been known [Vei71, Meg74]. To enable such a reduction, we need linear programming relaxations from the perspective of the signaling entity. Such a formulation is novel and non-trivial since a natural mathematical program from the perspective of the signaling entity is non-convex. This is because any program needs to encode the decision maker’s behavior. Handling this inherent non-convexity necessitates new structural insights (Lemma 4.3). Our LP formulation connects the signaling problem to classic stochastic optimization problems, such as stochastic knapsack [DGV08] and multi-armed bandits [GMS10]. This is in contrast to the more traditional and generic LP formulation for Bayesian persuasion [DX16, CDHW20] that directly encodes the optimal behavior of the decision maker, reducing the problem to multi-dimensional mechanism design [CDW12].

To build up to the general approximate majorization result, we first consider the simpler class of FULL REVELATION policies alluded to above that reveal all information to the decision maker. Though such policies cannot be approximately majorized in general (see Example 2.7), we show in Section 3.1 that when agent information is Bernoulli, such policies are without loss of generality, and are exactly majorized (hence best possible in terms of fairness). For arbitrary distributions, we show that there is always a 2-approximately majorized policy within the class of FULL REVELATION policies (see Section 3.2). These results showcase the network flow approach, which may be of independent interest in devising fair policies for other signaling problems.

1.1 Related Work

Bayesian Persuasion. Our selection model is a special case of *information design* (see [BM19, Dug17] for surveys) where an information mediator provides information to impact the behavior of one or more decision makers. This has also been termed *signaling* or *persuasion* in the literature. In the original model of *Bayesian persuasion* by [KG11], there is one agent called the *receiver* who receives additional information from a better-informed *sender*. Given the signal, the receiver computes its posterior over the state of nature and chooses an action to maximize its own utility. The sender can design the signals so that the receiver, acting in its own interest, maximizes some utility function the sender cares about. This problem has been widely studied in various contexts, such as price discrimination [BBM15, BMSW24], security games [XRDT15], regret minimization [BTXZ21], and other economic settings [CH14].

As mentioned earlier, computational approaches to the general Bayesian persuasion problem have been proposed, with an FPTAS for any given objective function [DKQ16, DX16]. However, our goal of majorization requires simultaneous optimization for many objectives, and here, even showing an existence result requires new ideas. Our techniques, starting with our LP relaxations from the perspective of the signaling entity, are entirely different. As a consequence, unlike [DX16], our methods work when the signaling scheme is independent across agents, while the computational results in [DX16] require correlation between the signals.

Majorization. The concept of majorized vectors is a classical one [Kar32, HLP34], and is equivalent to solutions that simultaneously maximize symmetric concave functions of the coordinates. In the context of resource allocation and routing problems, an approximate version of this concept was defined by [GMP05, GM06]; see also [KK06, CS19]. It was shown by [GMP05] that the best approximation factor is the solution of a linear program. However, the approximation factor is problem-dependent and can be linear in the number of dimensions. One major exception is the single-source multi-sink network flow problem, for which the elegant result of Veinott [Vei71] shows exact majorization. Our main contribution is showing a surprising connection between this work and the selection problem. We show that the approximation

factor for our problem is only logarithmic in the range of the quality scores, hence achieving stronger fairness properties than what the generic bounds would indicate.

We note that the work of [BMSW24] presents a constant-approximation to majorization for the special case of price discrimination. However, the state of nature in price discrimination is single dimensional (value of a single buyer), while it is multi-dimensional for our setting (joint quality distribution for n agents). Indeed, a constant-approximation is provably unachievable in our setting, and we need new technical ideas to show our approximation results. As far as we are aware, our work gives the first majorization results for persuasion problems with multi-dimensional type spaces.

Fair Selection. The meritocratic policy of selecting the candidate(s) with the highest evaluation scores is particularly natural and widely used. However, this policy can be at odds with fairness considerations. For instance, in a college admissions setting, standardized test outcomes can have correlations with demographic factors [DREM13], and these correlations can undesirably and significantly skew demographic outcomes. To mitigate this, one existing approach is to have a quota system setting aside a number of spots for candidates from certain demographic groups. A prominent example is the Rooney Rule, which among other instructions, requires that each National Football League (NFL) team must “employ a female or minority coach as an offensive assistant”.³ Different versions of the Rooney Rule have been seen in the corporate world, in governments, and in academia – those rules require employers to hire or at least interview candidates from prespecified demographic groups. Researchers have been studying the effectiveness of Rooney Rule variants in fair selection and hiring [KR18, CMV20].

Policies that consider demographic factors can fail when demographic information is unavailable, or when there is no consensus on which demographic groups should be protected. Interventions that explicitly consider demographic factors may also face societal and legal challenges. For instance, some practices have been deemed unlawful in U.S. Supreme Court decisions, such as racial quotas (Regents of the University of California v. Bakke (1978)) and more recently, race-based affirmative action (Students for Fair Admissions v. Harvard (2023)). Additionally, revealing demographic information can have adverse impact in hiring situations on, for example, women [GR00] and African Americans [BM04].

Departing from the above approaches, recent work explicitly models uncertainty in evaluation measures [EGGL20, MC21, GB21, SKJ21, SWZ⁺23, DKKS24] and proposes fair randomized selection rules under models of fairness and uncertainty over applicant qualities and attributes. Our work falls in this framework, but differs in the following way: We incorporate information revelation about applicant qualities, in order to (Bayesian-)persuade the meritocratic decision maker to act in line with fairness considerations.

Selection via Persuasion. Our model is similar to the selection models in [AK20, DTWZ24]. Unlike our model, in their model, there is no intermediary, and the agents are free to choose their signaling schemes. They show that the resulting game over the agents has an equilibrium under fairly general conditions. In contrast, we consider the setting with an intermediary that seeks fairness across agents; in other words, we assume agents can coordinate their signaling scheme for mutual benefit. Further, the models in [AK20, DTWZ24] assume the agent utilities are different from the social planner’s. In particular, agent utilities are binary – 1 for being selected and 0 otherwise – while we assume they are the same as their quality. Our setting models how the intermediary or social planner will perceive utility (as the quality), while in their setting, the utility is how it would have been perceived by selfish agents. We note however that our results are robust to the choice of utilities, and as we show in [Section 4.4](#), it is easy to modify our positive results to work as is in their utility model as well.

³<https://operations.nfl.com/inside-football-ops/inclusion/inclusive-hiring>, accessed June 28, 2024.

2 Our Model and Results

A social planner wishes to select one agent from a set A of n agents $\{1, 2, \dots, n\}$. Each agent i obtains value v_i if selected, where v_i is drawn from a distribution D_i . We assume that the distributions $\{D_i\}_{i=1}^n$ are mutually independent. In keeping with the Bayesian persuasion literature, we call the social planner the *receiver*. The receiver knows the distributions $\{D_i\}_{i=1}^n$, but not the values themselves, which are private information to the agents.

There is an information intermediary who knows the exact values $\{v_i\}_{i=1}^n$. Again following the Bayesian persuasion literature, we call the information intermediary the *sender*. The sender can reveal partial information about the values to the receiver via a *signaling policy*. (We will define signaling policies and signals in [Section 2.1](#).) The receiver hence receives a *signal* σ , and uses σ to update the prior value distributions $\{D_i\}_{i=1}^n$ to posterior distributions $\{D_i(\sigma)\}_{i=1}^n$.

2.1 Signaling Model

Independent Signaling Schemes. An *independent signaling scheme* ω works with a set Γ of signals, and we will often use $\sigma \in \Gamma$ to denote a signal. An independent signaling scheme in our model has two components: the *mapping rule* and the *selection rule*.

The *mapping rule* is a function that maps the values $\vec{v} = \{v_i\}_{i=1}^n$ to a distribution $g_{\vec{v}}$ over the signals Γ . The function is “independent” for each agent: There is a set Γ_i of signals for each agent i . The sender, after observing $v_i \sim D_i$, maps it to a distribution g_{iv} over signals in Γ_i . The sender then generates each signal $\sigma_i \sim g_{iv}$ and sends the set of generated signals $\sigma = \{\sigma_i\}_{i=1}^n$ to the receiver. The receiver computes a per-agent posterior $D_i(\sigma_i)$ (where $\{D_i(\sigma_i)\}_{i=1}^n$ are mutually independent). We can alternatively think that there is a separate sender for each agent that outputs a signal for that agent independent of the behaviors of other senders. This policy is known to the receiver.

After receiving the signal $\sigma \in \Gamma$, the receiver uses Bayes’ rule and its knowledge of the mapping policy to compute the independent posterior distributions $\{D_i(\sigma)\}_{i=1}^n$ over agent values. In our model, the receiver is a utilitarian welfare maximizer. Therefore, it will select an agent with the largest posterior mean – that is, an agent in the set $S(\sigma) := \arg \max_i \mathbb{E}[D_i(\sigma)]$. In the case of tie-breaking (i.e., when $|S(\sigma)| > 1$), we assume that the sender can tell the receiver who to select within the set $S(\sigma)$ via a *selection rule* that picks one agent in the set $S(\sigma)$ either deterministically or probabilistically. The receiver will follow the recommendation of the selection rule.

Signaling Policies. A *signaling policy* is a distribution Ω over independent signaling schemes. The receiver draws an independent signaling scheme $\omega \sim \Omega$ from this distribution and implements it. We use Ω to denote both the signaling policy and the distribution.⁴

2.2 Fairness via Approximate Majorization

Given a signaling policy Ω , we use $U_i(\Omega)$ to denote the expected utility that agent i receives, where the expectation is over the randomness in the signaling policy and the value distributions of the agents. This is formalized in [Definition 2.1](#).

Definition 2.1 (expected utility). The *expected utility* of agent i from Ω is

$$U_i(\Omega) = \mathbb{E}_{\omega \sim \Omega} \left[\sum_{\sigma \in \Gamma} q_{\omega}(\sigma) \cdot w_i(\sigma) \cdot \mu_i(\sigma) \right].$$

⁴One can also define more general versions of signaling policies that create signals with more correlation among the agents, but our definition aligns with our motivating practical applications.

In this formula, after $\omega \sim \Omega$ is drawn, $q_\omega(\sigma)$ denotes the probability (over agent values $\vec{v} \sim \prod_i D_i$) that the receiver observes the joint signal σ ; $w_i(\sigma)$ denotes the probability that agent i is selected when the signal is σ ; $\mu_i(\sigma) = \mathbb{E}[D_i(\sigma)]$ denotes the posterior mean value of agent i given signal σ .

The goal of the sender is to design a signaling policy Ω that is *fair* with respect to the expected utilities $\{U_i(\Omega)\}_{i=1}^n$. This is captured by the notion of α -majorization [GMP05, GM06], an extension to the classical mathematical notion of majorization [Kar32, HLP34].

Definition 2.2 (α -majorization). For $\alpha \geq 1$, a signaling policy Ω is called α -majorized⁵ if for any $k \in \{1, 2, \dots, n\}$ and any signaling policy Ω' , the sum of the k smallest utilities in $\{U_i(\Omega)\}_{i=1}^n$ is at least $1/\alpha$ times the sum of the k smallest utilities in $\{U_i(\Omega')\}_{i=1}^n$.

Sometimes, we use the phrase “ α -majorization within a class \mathcal{C} of signaling policies”. Its definition is to replace “any signaling policy Ω' ” to “any signaling policy $\Omega' \in \mathcal{C}$ ” in [Definition 2.2](#).

Denote $U(\Omega)$ as the vector $\{U_i(\Omega)\}_{i=1}^n$. The following result shows that approximate majorization is equivalent to simultaneously approximating all symmetric and concave welfare functions. It also holds when restricting to any class of signaling policies.

Proposition 2.3 (adapted from [GM06]). *The signaling policy Ω is α -majorized if and only if for every symmetric and concave function⁶ $f: \mathbb{R}_{\geq 0}^n \rightarrow \mathbb{R}_{\geq 0}$ (called a welfare function or fairness function) and any other signaling policy Ω' , it holds that*

$$f(U(\Omega)) \geq \frac{1}{\alpha} \cdot f(U(\Omega')).$$

Example 2.4. There are $n = 3$ agents. The value distributions D_1, D_2 of agents 1 and 2 are identical, and each of them takes value of 1 with probability 0.5 and takes value of 5 with probability 0.5. The value of agent 3 is deterministically 2. Consider the max-min fairness function: $f(U(\Omega)) = \min_{i=1}^n U_i(\Omega)$, so that the sender aims at maximizing the smallest expected utility among all the agents.

If the sender does not reveal any information about the agents, the receiver will select an agent from $\{1, 2\}$ because their expected values are both 3, which is larger than agent 3’s value of 2. Since agent 3 is never selected, the max-min welfare is 0.

Consider the following signaling policy Ω that implements an independent signaling scheme ω . The signal sets for each agent are: $\Gamma_1 = \{s_1, s'_1\}$, $\Gamma_2 = \{s_2, s'_2\}$, $\Gamma_3 = \{s_3\}$. The mapping rule of ω is given by the following probabilities:

$$\begin{aligned} \Pr[\sigma_1 = s_1 \mid v_1 = 1] &= 1; \quad \Pr[\sigma_1 = s_1 \mid v_1 = 5] = \frac{1}{3}; \quad \Pr[\sigma_1 = s'_1 \mid v_1 = 5] = \frac{2}{3}; \\ \Pr[\sigma_2 = s_2 \mid v_2 = 1] &= 1; \quad \Pr[\sigma_2 = s_2 \mid v_2 = 5] = \frac{1}{3}; \quad \Pr[\sigma_2 = s'_2 \mid v_2 = 5] = \frac{2}{3}; \\ \Pr[\sigma_3 = s_3] &= 1. \end{aligned}$$

The expected value of agent 1 in the receiver’s posterior upon receiving signal s_1 is

$$\begin{aligned} \mathbb{E}[D_1(s_1)] &= \mathbb{E}[v_1 \mid \sigma_1 = s_1] \\ &= 1 \cdot \Pr[v_1 = 1 \mid \sigma_1 = s_1] + 5 \cdot \Pr[v_1 = 5 \mid \sigma_1 = s_1] \\ &= 1 \cdot \Pr[\sigma_1 = s_1 \mid v_1 = 1] \cdot \frac{\Pr[v_1 = 1]}{\Pr[\sigma_1 = s_1]} + 5 \cdot \Pr[\sigma_1 = s_1 \mid v_1 = 5] \cdot \frac{\Pr[v_1 = 5]}{\Pr[\sigma_1 = s_1]} \\ &= 1 \times 0.75 + 5 \times 0.25 = 2. \end{aligned}$$

⁵The notion of “1-majorized” is also known as “least weakly supermajorized” in the literature [Tam95].

⁶Such a function will also be monotonically non-decreasing.

Note that $\mathbb{E}[D_1(s'_1)] = 5$. Similarly, we have $\mathbb{E}[D_2(s_2)] = 2$ and $\mathbb{E}[D_2(s'_2)] = 5$. Since D_3 is deterministic, we have $\mathbb{E}[D_3(s_3)] = 2$.

The selection rule of ω is to always break ties in favor of agent 3 and arbitrarily break ties between agent 1 and 2. Since agent 3 is selected when s_1 and s_2 are sent to the receiver, we have

$$U_3(\Omega) = \left(\frac{1}{3} \times 0.5 + 1 \times 0.5 \right)^2 \times 2 = \frac{8}{9}.$$

Since agent 1 is selected with value 5 as long as s'_1 and s_2 are sent, we have

$$U_1(\Omega) \geq \left(0.5 \times \frac{2}{3} \right) \times \left(\frac{1}{3} \times 0.5 + 1 \times 0.5 \right) \times 5 = \frac{10}{9} > \frac{8}{9} = U_3(\Omega).$$

Similarly, we have $U_2(\Omega) > U_3(\Omega)$. Therefore, $f(U(\Omega)) = \min\{U_1(\Omega), U_2(\Omega), U_3(\Omega)\} = 8/9$. We note that Ω is also the optimal signaling policy for the max-min welfare objective.

2.3 Our Results

FULL REVELATION Policy. If the sender aims at maximizing the utilitarian welfare $\sum_{i=1}^n U_i(\Omega)$ (which corresponds to the welfare function $f(\vec{x}) = \sum_{i=1}^n x_i$), then one optimal mapping rule is to fully reveal the values $\{v_i\}_{i=1}^n$. In this case, the sender's goal perfectly aligns with the receiver's selection behavior, obtaining the best possible utilitarian social welfare of $\mathbb{E}_{\vec{v} \sim \prod_i D_i} [\max_{i=1}^n v_i]$. We term this mapping rule **FULL REVELATION**. It is easy to check that when $\{D_i\}_{i=1}^n$ are i.i.d., **FULL REVELATION** paired with a symmetric selection rule is 1-majorized. Also, we call an independent signaling scheme **FULL REVELATION** if its mapping rule is **FULL REVELATION**.

Our first result shows that when each D_i is $\text{Bernoulli}(\mu_i)$ (which takes value 1 with probability w_i and value 0 otherwise), there is a 1-majorized **FULL REVELATION** independent signaling scheme – it simultaneously optimizes all welfare functions, not just the utilitarian welfare. Though this result is for specific value distributions, its main idea of reducing to network flow will be useful for an approximate majorization result in [Section 4](#) for more general distributions. The selection policy can be computed in polynomial time. (Here and later when we mention “polynomial time” results, we assume that the value distributions are discrete.)

Theorem 2.5 (Proved in [Section 3.1](#)). *When each D_i is $\text{Bernoulli}(\mu_i)$, **FULL REVELATION** is 1-majorized when paired with a certain polynomial-time-computable selection policy.*

To prove the above theorem, we reduce the problem to network flow and apply a seminal majorization result for flows from [\[Meg74\]](#) (see [Theorem 3.3](#)). A similar technique shows that even for general distributions, when restricted to the class of **FULL REVELATION** policies, there is a 2-majorized policy.

Theorem 2.6 (Proved in [Section 3.2](#)). *For general distributions $\{D_i\}_{i=1}^n$, **FULL REVELATION**, when paired with a certain polynomial-time-computable selection policy, is 2-majorized within the space of **FULL REVELATION** policies.*

Note that [Theorem 2.6](#) shows majorization *within* the class of **FULL REVELATION** policies. In contrast, the example below shows that if we consider the space of *all* signaling schemes, **FULL REVELATION** policies cannot achieve a good approximate majorization. This impossibility result motivates us to study other classes of signaling policies.

Example 2.7. The distribution D_1 is deterministically 2, and each of the distributions D_2, \dots, D_n is 1 with probability 1/2 and 3 with probability 1/2. Consider the welfare function of max-min fairness

where $f(U(\Omega)) = \min_i U_i(\Omega)$. In FULL REVELATION, $U_1 = 2/2^{n-1}$, and therefore the max-min welfare for FULL REVELATION is at most $2/2^{n-1}$. (To calculate the expected utilities of other agents, we have $\mathbb{E}[\max_{i=2}^n D_i] = 3 \cdot (1 - 1/2^{n-1}) + 1 \cdot (1/2^{n-1}) \approx 3$ when n is large. Therefore, for a symmetric selection policy, $U_2 = \dots = U_n \approx 3/n$.) In contrast, if the policy reveals nothing, all agents have posterior mean 2. Assuming the selection policy assigns uniformly at random, $U_1 = U_2 = \dots = U_n = 2/n$. Therefore, FULL REVELATION is no better than a $(2^{n-1}/n)$ -approximation to the max-min fairness objective.

Approximate Majorization. We now consider the question of majorization among *all* signaling policies. For the rest of our results, we assume that each D_i is supported on $[1, V]$. (This interval is equivalent to any $[v_{\min}, v_{\max}]$ with $V = v_{\max}/v_{\min}$ via scaling.) Our results characterize the approximation to majorization as functions of V . The main hurdle with formulating a mathematical program for finding such majorized signaling policies is the non-linear and non-convex interaction between variables encoding mapping and those encoding selection.

We overcome this challenge in two steps: First, we split the utility of a fixed policy into buckets. Second, we efficiently find a mapping that ensures a large utility in some bucket by solving a simple LP with constraints on posterior means, enabling a reduction to the network flow approach from before. For the second step, we consider a specific type of mapping rules which we call SINGLE MEAN. Such a mapping rule picks a common range $[\mu, \hat{\mu}]$ for all agents and finds a mapping for each agent that maximizes the probability that the mean of the resulting signal lies in the range. We will choose $\hat{\mu} = (1 + \varepsilon)\mu$ for some constant $\varepsilon > 0$. The signaling policies we consider randomize over SINGLE MEAN mapping rules for different μ , and we will describe the corresponding selection rules later. Note that the mapping rule for any agent does not depend on other agents, so that the mapping can be constructed in a distributed way.

We show that this class of policies suffices to achieve a non-trivial positive approximation result for majorization. Our result is *bicriteria*, in the following sense. Let Z_k denote the maximum achievable value of the sum of the smallest k utilities. For given $\varepsilon > 0$, we say that a policy Ω is a bicriteria α -approximation if the following hold.

- We allow Ω to use a $(1 + \varepsilon)$ -approximate welfare maximizer⁷ – that is, for any signal σ , the selection rule can select any agent i with $\mathbb{E}[D_i(\sigma)] \geq \max_j \mathbb{E}[D_j(\sigma)]/(1 + \varepsilon)$.
- For each k , the sum of the smallest k utilities in Ω is at least Z_k/α . Note that the quantity Z_k assumes the receiver is an exact welfare maximizer.

Theorem 2.8 (Proved in [Section 4](#)). *For any constant $\varepsilon > 0$, there is a bicriteria $O((\log V)/\varepsilon)$ -majorized polynomial-time-computable signaling policy.*

Remark. The generic result of [\[GM06\]](#) can be applied to our setting and give an $O\left(\min\left\{n, \log \frac{P_{\max}}{n P_{\min}}\right\}\right)$ -majorized solution, where $P_{\min} = \max_{\Omega} \min_i U_i(\Omega)$ and $P_{\max} = \max_{\Omega} \sum_i U_i(\Omega)$. In our setting, the former quantity could be exponentially (in n) smaller than the latter, leading to the approximation ratio having linear dependence on n . In contrast, our main result has no dependence on n and only logarithmic dependence on V , the range of values.

At a technical level, [Lemma 4.3](#) characterizes the structure about the class of SINGLE MEAN mappings, and enables us to pre-compute the mapping variables. We can therefore decouple the mapping variables from the selection variables, and reduce the non-convex mathematical program to a network flow problem. Finally, we use the majorization result in [Theorem 3.3](#) to complete the proof.

As an immediate corollary, this shows a polynomial-time $O((\log V)/\varepsilon)$ bicriteria approximation algorithm for any given fairness function f when agents generate mappings in a distributed fashion, which is

⁷We note that polynomial-time-computability results of [\[DX16\]](#) also assume an approximately optimal receiver.

desirable for the applications of fair selection. We note that the FPTAS in [DX16] assumes that the sender can use all agents' values to generate a common signal.

In Section 4.4, we extend the utility model to be asymmetric between the agents and the receiver, capturing the setting in [AK20]. In this model, the signals and the receiver's action depend on the quality of the agents, while an agent gets utility 1 if selected and 0 otherwise. We show that Theorem 2.8 extends as is, and shows approximate majorization even in this setting.

We complement the previous theorem with a lower bound showing a dependence on V in the approximation factor is unavoidable, even if the receiver is an approximate welfare maximizer.

Theorem 2.9 (Proved in Section 5). *No signaling policy can be $(\log \log V)/3$ -majorized, even if the selection rule can arbitrarily select any agent (i.e., without the constraint that the selected agent has near-optimal expected value.).*

3 Warmup: The FULL REVELATION Signaling Policy

In this section, we consider the FULL REVELATION signaling policy, in which the sender reveals all the values $\vec{v} = \{v_i\}_{i=1}^n$ to the receiver. The sender can still design the selection rule based on the fairness objective. In this case, the selection rule specifies which agent – among the ones that have the highest value – the receiver should select.

We propose a formulation of the problem based on network flow, and en route develop techniques that will also be useful in Section 4. In this section, using these techniques, we show that FULL REVELATION can be 1-majorized when the value distributions are Bernoulli (but can be heterogeneous). We then show that for general value distributions, when restricted to the space of FULL REVELATION signaling policies, there is a 2-majorized, polynomial-time-computable signaling policy.

3.1 1-Majorization for Bernoulli Distributions: Proof of Theorem 2.5

Consider the Bernoulli setting in which each value distribution D_i is $\text{Bernoulli}(\mu_i)$ (which takes value 1 with probability μ_i and takes value 0 otherwise). We first show in Lemma 3.1 that any signaling policy can be converted to using a FULL REVELATION signaling policy without decreasing the utility of any agent. This reduction allows us to only care about FULL REVELATION policies. We subsequently map each FULL REVELATION policy to a network flow instance, thus enabling us to apply the majorization result for network flows [Vei71].

To prove Lemma 3.1, we construct a new FULL REVELATION signaling policy Ω' that reveals the full value vector of all the agents to the receiver. We carefully set the selection policy of Ω' so that each agent's winning probability conditioned on their value being 1 does not decrease in Ω' compared to Ω .

Lemma 3.1. *When $D_i = \text{Bernoulli}(\mu_i)$, for any signaling policy Ω , there exists a FULL REVELATION signaling policy Ω' in which the utility of each agent is at least their utility in Ω .*

Proof. Given any signaling policy Ω , let Γ be the set of all possible signals that can be sent to the receiver in Ω . We will construct a FULL REVELATION signaling policy Ω' that preserves the utilities of all agents.

For each $\sigma \in \Gamma$, recall that $D_i(\sigma)$ is the receiver's posterior Bernoulli distribution over agent i 's value. Denote by $q_i(\sigma)$ the conditional probability that agent i is selected when σ is sent to the receiver.

In Ω , suppose the sender observes the value vector of agents as \vec{v}_0 and sends signal $\sigma \in \Gamma$. Then, in Ω' , the sender reveals \vec{v}_0 to the receiver. Suppose that $S(\vec{v}_0)$ is the set of agents with value 1 in \vec{v}_0 . In Ω' , the selection policy of the receiver chooses $i \in S(\vec{v}_0)$ with probability $q_i(\sigma)$. If $\sum_{i \in S(\vec{v}_0)} q_i(\sigma) < 1$, the selection probabilities for $S(\vec{v}_0)$ are arbitrarily increased so that they sum to 1. Denote v_0^i as the i^{th}

component of \vec{v}_0 . Noting that an agent only gains utility when their value is 1, in the new signaling policy Ω' , the expected utility of agent i is

$$\begin{aligned}
U_i(\Omega') &\geq \sum_{\vec{v}_0: v_0^i=1} \Pr[\vec{v} = \vec{v}_0] \cdot \left(\sum_{\sigma \in \Gamma} \Pr[\sigma \text{ is sent in } \Omega \mid \vec{v} = \vec{v}_0] \cdot q_i(\sigma) \right) \\
&= \sum_{\sigma \in \Gamma} \left(\sum_{\vec{v}_0: v_0^i=1} \Pr[\sigma \text{ is sent in } \Omega, \vec{v} = \vec{v}_0] \cdot q_i(\sigma) \right) \\
&= \sum_{\sigma \in \Gamma} \left(\sum_{\vec{v}_0: v_0^i=1} \Pr[\vec{v} = \vec{v}_0 \mid \sigma \text{ is sent in } \Omega] \right) \cdot \Pr[\sigma \text{ is sent in } \Omega] \cdot q_i(\sigma) \\
&= \sum_{\sigma \in \Gamma} \Pr[v_i = 1 \mid \sigma \text{ is sent in } \Omega] \cdot \Pr[\sigma \text{ is sent in } \Omega] \cdot q_i(\sigma).
\end{aligned}$$

In Ω , conditioned on σ being sent, the expected utility of agent i is $q_i(\sigma) \cdot \mathbb{E}[D_i(\sigma)]$. Since $D_i(\sigma)$ is Bernoulli, we have $\mathbb{E}[D_i(\sigma)] = \Pr[v_i = 1 \mid \sigma \text{ is sent in } \Omega]$. Therefore,

$$U_i(\Omega) = \sum_{\sigma \in \Gamma} \Pr[v_i = 1 \mid \sigma \text{ is sent in } \Omega] \cdot \Pr[\sigma \text{ is sent in } \Omega] \cdot q_i(\sigma) \leq U_i(\Omega').$$

This completes the proof. \square

Reduction to Network Flow. We now construct the network flow instance from the FULL REVELATION policies. We first build a bipartite directed flow graph G as follows:

- There is a single source s . Place a set R of n sinks t_1, \dots, t_n where t_i represents the agent i .
- Place a set L of 2^n nodes where each subset $T \subseteq [n]$ corresponds to a node s_T .
- Add an edge (s, s_T) with capacity $\mu_T = \prod_{i \in T} \mu_i \cdot \prod_{i \notin T} (1 - \mu_i)$. For each agent $i \in T$, add an edge (s_T, t_i) with capacity $+\infty$.

We illustrate this construction on an example with $n = 2$ agents in Figure 1. Their value distributions are $D_1 = \text{Bernoulli}(0.3)$ and $D_2 = \text{Bernoulli}(0.6)$. Since there are two agents, we have 4 nodes in L , denoted by $s_{\emptyset}, s_{\{1\}}, s_{\{2\}}$ and $s_{\{1,2\}}$. By the construction above, we have the capacities on edges from s to these four nodes are $(1 - 0.3) \times (1 - 0.6) = 0.28$, $0.3 \times (1 - 0.6) = 0.12$, $(1 - 0.3) \times 0.6 = 0.42$ and $0.3 \cdot 0.6 = 0.18$, respectively.

The intuition for the construction and Lemma 3.2 is the following. Each node in L corresponds to a possible value vector of all agents, or equivalently, the set of agents whose values are 1. The flow value from the source to such node is capped by the probability that this value vector is realized. The flow value from a node in L to an agent node equals the probability of the event that the vector corresponding to L is realized and this agent is selected.

Lemma 3.2. *The set of feasible s - t maximum flows in G corresponds exactly to the set of FULL REVELATION signaling policies. For any policy Ω , the utility $U_i(\Omega)$ of agent i is equal to the total flow f_i to the sink t_i in G in the corresponding flow instance.*

Proof. We first prove that any signaling policy Ω can be converted to a maximum flow. Fix any set $T \subseteq [n]$, and consider the event in which the agents with value 1 exactly form the set T . Let z_T be the probability

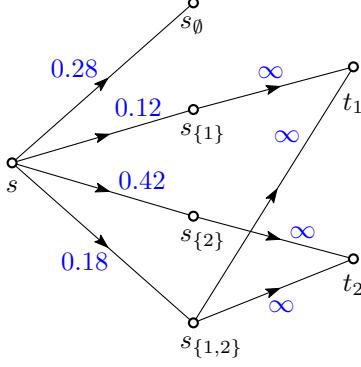


Figure 1: The network flow instance for an example with $n = 2$ agents, with $D_1 = \text{Bernoulli}(0.3)$ and $D_2 = \text{Bernoulli}(0.6)$. The number in blue next to each directed edge indicates the capacity of the edge.

that conditioned on this event, the agent $i \in T$ is selected in Ω . We set flow f_{Ti} on the edge (s_T, t_i) to be $\mu_T \cdot z_{Ti}$. Note that $U_i(\Omega) = \sum_{T:i \in T} \mu_T \cdot z_{Ti} = \sum_{T:i \in T} f_{Ti} = f_i$. The total flow is $\sum_{T \neq \emptyset} \mu_T$, which means that it is a maximum flow.

For the other direction, given a maximum flow, let T denote the set of agents with value 1. Notice that when $T \neq \emptyset$, the flow from s to s_T must be equal to its capacity. We construct a FULL REVELATION signaling policy Ω by requiring that conditioned on T being the set of agents with value 1, the receiver selects $i \in T$ with probability f_{Ti}/μ_T , where f_{Ti} is the flow on the edge (s_T, t_i) . This is clearly a feasible signaling policy with $U_i(\Omega) = \sum_{T:i \in T} \mu_T \cdot (f_{Ti}/\mu_T) = f_i$. \square

The following theorem (implicit in [Vei71]) shows the existence of a 1-majorized flow, and therefore shows the existence of a 1-majorized FULL REVELATION signaling policy. We include a short proof for completeness.

Theorem 3.3 (1-Majorized Flows, Implicit in [Vei71]). *Given any capacitated network with a single source and multiple sinks t_1, \dots, t_n , let $f = \{f_i\}_{i=1}^n$ be the flows entering the respective sinks in a feasible flow. Then there is a flow f^* that is 1-majorized among these feasible flows.*

Proof. By [Meg74, Lemma 4.1], the set of feasible $\{f_i\}_{i=1}^n$ defines a polymatroid – that is, for all $T \subseteq [n]$, it holds that $\sum_{i \in T} f_i \leq g(T)$ and $\{f_i\}_{i=1}^n \geq 0$, where g is a non-negative, monotone, submodular function. By [Tam95, Theorem 3.2] (see also [DR89]), any polymatroid has a 1-majorized point. \square

Polynomial-Time Algorithm. Though the graph G constructed above has exponential size, it follows from [Meg74] that a flow $\{f_i\}_{i=1}^n$ is feasible if and only if it is feasible for the following polymatroid \mathcal{P} :

$$\sum_{i \in T} f_i \leq \sum_{J: J \cap T \neq \emptyset} \mu_J = 1 - \prod_{i \in T} (1 - \mu_i), \quad \forall T \subseteq [n];$$

$$f_i \in [0, 1], \quad \forall i \in [n].$$

Maximizing any strictly concave, symmetric, separable function over this polymatroid now yields the solution in polynomial time [Tam95, Vei71]. (See [GM06] for an LP based approach.) This can be solved to an arbitrary approximation in polynomial time.

The final solution f^* for G lies in the convex hull of the vertices of \mathcal{P} . By Carathéodory's Theorem, this point can be written as the convex combination of at most $n+1$ vertices on the convex hull. By applying an ellipsoid method, we can find such a decomposition in polynomial time. Denote the decomposed vertices and their weights by $\{(w_k, \eta_k)\}_{k=1}^{n+1}$, where $\sum_{k=1}^{n+1} \eta_k = 1$.

By the property of polymatroids, each vertex u_t corresponds to a series of “tight subsets” $\emptyset = T_0 \subsetneq T_1 \subsetneq T_2 \subsetneq \cdots \subsetneq T_n = [n]$, where each adjacent pair of sets (T_{k-1}, T_k) only differ by a single agent i_k . By a “tight subset” T we mean that the inequality $\sum_{i \in T} f_i \leq 1 - \prod_{i \in T} (1 - \mu_i)$ is tight at this vertex. Since f_i is the utility of agent i for this vertex solution, as long as at least one agent in T_k has value 1, some agent in this set must be finally selected. This yields the following ranking scheme for w_t :

1. The sender reveals all indices $\{i_k\}_{k=1}^n$.
2. The receiver selects the first agent i_k in the ranking whose value is 1.

The signaling scheme corresponding to f^* is a randomization over these ranking schemes: With probability η_k , run the ranking mechanism corresponding to the vertex w_k . Since the utilities of the agents are exactly $\{f_i^*\}_{i=1}^n$, this is 1-majorized. This completes the proof of [Theorem 2.5](#).

3.2 Majorization among FULL INFORMATION Schemes: Proof of [Theorem 2.6](#)

We now show that if we restrict to the space of FULL INFORMATION mapping policies, there is a selection policy that is 2-majorized and can be computed in polynomial time. Our technique again reduces to network flow, albeit for a relaxation of the problem. We will use a similar relaxation in [Section 4](#) for the general problem.

We assume that there are m possible values in total that support the value distributions $\{D_i\}_{i=1}^n$. In the remainder of this section, we use the notation $v_1 < v_2 < \cdots < v_m$ to denote these m possible values. Let $d_{ij} = \Pr[D_i = v_j]$. Let Z_j be the event that the largest value is at most v_j . We further define

$$z_j := \Pr[Z_j] = \prod_{i=1}^n \left(\sum_{j' \leq j} d_{ij'} \right)$$

and

$$p_{ij} := \Pr[D_i = v_j \mid Z_j] = \frac{d_{ij}}{\sum_{j' \leq j} d_{ij'}}.$$

Network Flow Instance. We will show a network flow instance so that any FULL REVELATION policy Ω can be relaxed into a feasible flow for this instance. Given Ω , let x_{ij} be the probability that agent i has value v_j and is selected conditioned on the event Z_j . For the policy Ω , the variables $\{x_{ij}\}$ and expected utilities $\{u_i\}_{i=1}^n$ satisfy the following constraints.

$$\sum_{j=1}^m z_j \cdot v_j \cdot x_{ij} = u_i, \quad \forall i \in [n]; \tag{1}$$

$$\sum_{i=1}^n x_{ij} \leq 1, \quad \forall j \in [m]; \tag{2}$$

$$0 \leq x_{ij} \leq p_{ij}, \quad \forall i \in [n], j \in [m]. \tag{3}$$

Conditioned on the event Z_j , [Eq. \(2\)](#) says at most one agent with value v_j is selected, and [Eq. \(3\)](#) says the probability that agent i has value v_j and is selected is at most p_{ij} , the probability that its value is v_j .

Let $y_{ij} = z_j \cdot v_j \cdot x_{ij}$, and $b_j = z_j \cdot v_j$. We can rewrite the above program as:

$$\begin{aligned} \sum_{j=1}^m y_{ij} &= u_i, & \forall i \in [n]; \\ \sum_{i=1}^n y_{ij} &\leq b_j, & \forall j \in [m]; \\ 0 \leq y_{ij} &\leq p_{ij} \cdot b_j, & \forall i \in [n], j \in [m]. \end{aligned}$$

Construct a network flow graph G with a source s , and one sink t_i for each agent i . There are m nodes, one for each v_j . There is a directed edge (s, v_j) with capacity b_j , and a directed edge (v_j, t_i) with capacity $p_{ij} \cdot b_j$. The flow value on the edge (v_j, t_i) is y_{ij} , and sink t_i receives flow u_i . This shows the network flow instance such that any Ω relaxes to a feasible flow for this instance, and the structure of the network is illustrated in Fig. 2.

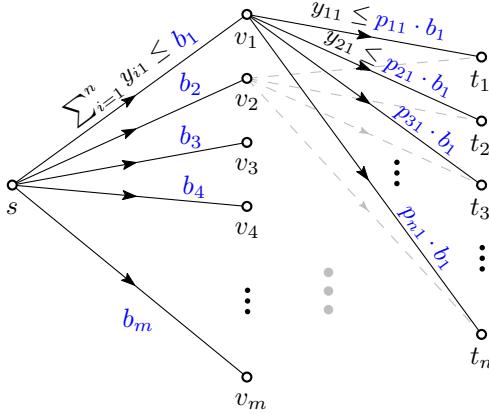


Figure 2: Illustration of the network flow instance. The text in blue denotes the capacity of each edge, while the text in black denotes the feasible flow value on each edge.

Using [Theorem 3.3](#), there is a 1-majorized flow in G , which translates to a 1-majorized utility vector \vec{u} . Further, this solution can be computed in polynomial time [[Tam95](#), [Vei71](#), [GM06](#)]. Denote the majorized solution (in the original variables) as $\{x_{ij}\}$ and utilities as $\{u_i\}$.

Signaling Policy. We now convert the 1-majorized solution to the above flow instance into a simple 2-majorized signaling policy. This is inspired by the algorithm for stochastic knapsack [DGV08].

- The sender sends all realized values to the receiver.
- Suppose that the largest realized value is v_j . Let S denote the set of agents i with value v_j .
- Sort the agents in S in a uniformly random order, and consider the agents in this order.
- In this order, when we reach each agent i , with probability x_{ij}/p_{ij} , select i and stop.
- If no agent in S is selected above, then select an arbitrary agent in S .

Analysis. Consider any agent i . Conditioned on the event Z_j (i.e., the largest value being at most v_j), the probability that another agent $i' \in S$ is selected before agent i is reached is at most

$$\frac{1}{2} \cdot \frac{x_{i'j}}{p_{i'j}} \cdot p_{i'j} = \frac{x_{i'j}}{2},$$

where the factor of $1/2$ comes from the fact that the order is uniformly random. Since $\sum_{i'} x_{i'j}/2 \leq 1/2$, the probability that no agent before i is selected is at least $1/2$. Now, for agent i with value v_j to be selected, the following should happen:

- The event Z_j , which happens with probability z_j ;
- Agent i has value v_j conditioned on Z_j , which happens with probability p_{ij} ;
- No agent before i is selected conditioned on all previous events, which happens with probability at least $1/2$; and
- Agent i with value v_j is selected conditioned on all previous events, which happens with probability x_{ij}/p_{ij} .

Therefore, the unconditional probability that agent i with value v_j is selected is

$$\Pr[\text{agent } i \text{ with value } v_j \text{ is selected}] \geq z_j \cdot \frac{1}{2} \cdot p_{ij} \cdot \frac{x_{ij}}{p_{ij}} = z_j \cdot \frac{x_{ij}}{2}.$$

Therefore, the expected utility of agent i in this scheme is at least:

$$\sum_{j=1}^m \Pr[\text{agent } i \text{ with value } v_j \text{ is selected}] \cdot v_j \geq \sum_{j=1}^m z_j \cdot \frac{x_{ij}}{2} \cdot v_j = \frac{u_i}{2}. \quad (4)$$

This shows a 2-majorized signaling scheme running in polynomial time, proving [Theorem 2.6](#).

4 Approximate Majorization: Proof of [Theorem 2.8](#)

We now present our main technical result – approximate majorization for general signaling policies. The key hurdle with using the approach in [Sections 3.1](#) and [3.2](#) is that unlike the FULL REVELATION setting where the mapping rule is fixed, any mathematical program for general policies needs to encode both mapping and selection variables, and their interaction is non-linear. To deal with the challenge, we will perform a *projection* of any signaling policy onto a specific class of SINGLE MEAN policies, for which the network flow majorization approach in the previous sections can be employed.

4.1 SINGLE MEAN Projections

Given a small constant $\varepsilon > 0$, let $\eta = 1 + \varepsilon$. Without loss of generality, assume V is a power of η , and divide the interval $[1, V]$ into buckets $I_1 = [1, \eta], I_2 = [\eta, \eta^2], \dots, I_K = [V/\eta, V]$, where $K \approx (\log V)/\varepsilon$. For any policy Ω^* , let u_i denote agent i 's utility, that is, $u_i = U_i(\Omega^*)$. We can express u_i as the sum of contributions from different buckets as $u_i = \sum_{k=1}^K c_{ik}$, where c_{ik} is the contribution to the utility u_i from those signals σ for which the posterior mean satisfies $\mu_i(\sigma) := \mathbb{E}[D_i(\sigma_i)] \in I_k$.

We now define SINGLE MEAN projections that assume that the only contribution to utility for agent i happens when the signal σ received has posterior mean $\mu_i(\sigma) \in I_k$ and agent i is selected.

Definition 4.1 (SINGLE MEAN projection). Given an interval I_k and signaling policy Ω , let A_{ik} denote the event where both (I) agent i is selected and (II) $\mathbb{E}[D_i(\sigma)] \in I_k$. The SINGLE MEAN projection Ω_k of Ω onto I_k only counts utilities from the event $\bigcup_{i=1}^n A_{ik}$, and sets the remaining utilities to 0. In other words, if $c_{ik}(\Omega) = \Pr[A_{ik}] \cdot \mathbb{E}[\mu_i(\sigma) | A_{ik}]$, then $U_i(\Omega_k) = c_{ik}$ for all agents i .

Consider now the policy that picks a number $k \in \{1, 2, \dots, K\}$ uniformly at random and executes the SINGLE MEAN projection Ω_k of Ω^* for this interval. The expected utility of agent i is exactly $\sum_{k=1}^K c_{ik}/K = u_i/K$. We overload the notation and denote the new utilities as $\{u_i\}_{i=1}^n$.

4.2 Structure of SINGLE MEAN Projections

We now show that SINGLE MEAN projections can be modified to have a surprisingly simple structure while not decreasing the utilities $\{c_{ik}\}$. Let the interval under consideration be $I_k = [m, \eta \cdot m]$ (or $[m, \eta \cdot m]$ if $k = K$) and let the SINGLE MEAN projection of Ω^* be Ω_k . Losing an additional factor η in the approximation factor, we can assume that if $\mu_i(\sigma) \in I_k$, then the expected utility of agent i is always m if i is selected.

Note that if $\mu_i(\sigma) \in I_k$ and agent i is selected, then in this scenario, $\max_j \mu_j(\sigma) \in [m, \eta \cdot m]$. Denote the interval $[m, \eta \cdot m]$ as \hat{I} , and $\eta \cdot m$ as \hat{m} . Assume now that when all agents have posterior mean at most \hat{m} , any agent whose posterior mean is in \hat{I} can be selected, and this choice yields utility m to that agent. Such a policy assumes an η -approximate utilitarian-welfare-maximizing receiver, and hence our final result will be bicriteria.

Clearly, whenever Ω_k selects an agent with posterior mean in I_k , all agents must have posterior means at most \hat{m} , so that the new policy can also be made to make the same selection. If we pretend the utility yielded to the agent is m , this is still within factor η of the actual utility – the modified policy has utilities $\hat{c}_{ik} \geq c_{ik}/\eta$.

Mathematical Program. Focus on some k , and denote the relaxed SINGLE MEAN projection of Ω^* as Ω_k^* . This policy can be written as a distribution over independent signaling schemes as

$$\Omega_k^* = \sum_{\ell \geq 1} \rho_\ell \cdot \tau_\ell,$$

where each independent signaling scheme τ_ℓ is an independent mapping of the agents to signals, along with an associated selection rule, and $\sum_{\ell \geq 1} \rho_\ell = 1$. For any independent signaling scheme τ in the above summation, let $u_i(\tau)$ denote the utility of agent i in this scheme, so that

$$U_i(\Omega_k^*) = \sum_{\ell} \rho_\ell \cdot u_i(\tau_\ell).$$

Note that given the SINGLE MEAN structure, we only count the utility when the posterior mean of the selected agent lies in \hat{I} , in which case the contribution is assumed to be m .

Given signaling scheme τ , let $B(\tau)$ denote the event that every agent j has $\mu_j(\sigma) \leq \hat{m}$, where $\sigma \sim \tau$ is an independent signal. Now define the following four probabilities for the signaling scheme τ :

$$\begin{aligned} Q_i(\tau) &:= \Pr_{\sigma \sim \tau} [\mu_i(\sigma) \leq \hat{m}]; \\ Q(\tau) &:= \Pr[B(\tau)] = \prod_{i=1}^n Q_i(\tau); \\ p_i(\tau) &:= \Pr_{\sigma \sim \tau} [\mu_i(\sigma) \in \hat{I} | B(\tau)]; \\ x_i(\tau) &:= \Pr_{\sigma \sim \tau} [\text{agent } i \text{ is selected and } \mu_i(\sigma) \in \hat{I} | B(\tau)]. \end{aligned}$$

$Q(\tau) = \prod_{i=1}^n Q_i(\tau)$ comes from the assumption that in the independent signaling scheme τ , the mapping rule of each agent is independent of all other agents. The utilities in policy τ clearly satisfy:

$$u_i(\tau) = m \cdot Q(\tau) \cdot x_i(\tau), \quad \forall i \in [n]. \quad (5)$$

The event $\mu_i(\sigma) \in \hat{I}$ conditioned on $B(\tau)$ follows $\text{Bernoulli}(1, p_i(\tau))$, so that analogous to [Section 3.1](#), the following polymatroid exactly captures the $\{x_i(\tau)\}_{i=1}^n$ in any feasible policy.

$$\sum_{i \in S} x_i(\tau) \leq 1 - \prod_{i \in S} (1 - p_i(\tau)), \quad \forall S \subseteq [n]. \quad (6)$$

Structural Lemma. The issue now is that the quantities $Q(\tau)$ and $p_i(\tau)$ depend on the policy τ , so that we cannot yet reduce it to network flow in a way analogous to [Section 3.1](#). Our key insight is that this dependence can be removed. Towards this end, we define a *maximal mapping* as follows.

Definition 4.2 (Maximal Mapping). For an interval \hat{I} , a *maximal mapping* is a mapping rule ω from agent values to signals $\{\sigma\}$ such that $\Pr_{\sigma \sim \omega}[\mu_i(\sigma) \in \hat{I}]$ is maximized for each agent i .

We first show the following structural lemma, and subsequently show how to compute the maximal mapping efficiently.

Lemma 4.3. *Given any independent signaling scheme τ above, there is an independent signaling scheme τ' that (I) uses a maximal mapping, (II) satisfies [Eqs. \(5\)](#) and [\(6\)](#), and (III) gives each agent i a utility of at least $u_i(\tau)$ from the signals in which their posterior mean is inside \hat{I} , where the contribution to the utility is assumed to be m .*

Proof. Consider τ and the corresponding mapping rule for agent i . Let S_i^1 be the set of signals σ with the posterior mean $\mu_i(\sigma) < m$; S_i^2 be those with $\mu_i(\sigma) \in \hat{I}$; and S_i^3 those with $\mu_i(\sigma) > \hat{m}$. Create three signals $\sigma_i^1, \sigma_i^2, \sigma_i^3$. These are sent whenever a signal in S_i^1, S_i^2, S_i^3 is sent, respectively. Note that $\mu_i(\sigma_i^1) = m_i^1 < m$; $\mu_i(\sigma_i^2) = m_i^2 \in \hat{I}$; and $\mu_i(\sigma_i^3) = m_i^3 > \hat{m}$. Further, $Q_i = \sum_{\sigma \in S_i^1 \cup S_i^2} \Pr[\sigma] = \Pr[\sigma_i^2] + \Pr[\sigma_i^1]$ remains unchanged and similarly p_i remains unchanged. This preserves both the constraints and the utilities.

For agent i , let α_i satisfy $\alpha_i m_i^1 + (1 - \alpha_i) m_i^3 = m_i^2$, and define $s_i^1 := \Pr[\sigma_i^1]$ and $s_i^3 := \Pr[\sigma_i^3]$. Create a new signal σ_i^4 , and do the following things:

- If $\beta = \frac{s_i^1(1-\alpha_i)}{s_i^3\alpha_i} \leq 1$, then whenever signal σ_i^1 was sent, σ_i^4 is sent instead, and whenever σ_i^3 was sent, σ_i^4 is sent instead with probability β and σ_i^3 is sent with probability $1 - \beta$.
- If $\beta > 1$, then whenever signal σ_i^3 was sent, σ_i^4 is sent instead, and whenever σ_i^1 was sent, σ_i^4 is sent instead with probability $1/\beta$ and σ_i^1 sent with probability $1 - 1/\beta$.

In either case, we note that $\mu_i(\sigma_i^4) = \alpha_i m_i^1 + (1 - \alpha_i) m_i^3 = m_i^2 \in \hat{I}$. The sender can then send signal σ_i^2 whenever σ_i^4 is sent, noting that this preserves the condition that $\mathbb{E}[\mu_i(\sigma_i^2)] \in \hat{I}$.

Note that in either case, $\Pr[\sigma_i^3]$ does not increase, and hence $Q_i = \Pr[\sigma_i^1] + \Pr[\sigma_i^2]$ does not decrease. Since the increase in $\Pr[\sigma_i^2]$ is exactly the decrease in the sum $\Pr[\sigma_i^1] + \Pr[\sigma_i^3]$, we know that p_i does not decrease. Therefore, the solution $\{x_i(\tau)\}_{i=1}^n$ remains feasible for [Eq. \(6\)](#) with the variables $\{p_i\}_{i=1}^n$ for the new mapping, and the right-hand side in [Eq. \(5\)](#) does not decrease as well, since $Q = \prod_{i=1}^n Q_i$ does not decrease. Therefore, the utility $u_i(\tau)$ does not decrease. [Figure 3](#) illustrates the reduction process for the two different cases discussed above.

For any agent i , the above process yields two signals σ_i^a, σ_i^b , where $\mu_i(\sigma_i^a) \in \hat{I}$. Recall that $\mu_i = \mathbb{E}[D_i]$ is the prior mean, and is a convex combination of $\mu_i(\sigma_i^a)$ and $\mu_i(\sigma_i^b)$. There are three cases now.

1. If $\mu_i \in \hat{I}$, then not revealing anything yields $Q_i = p_i = 1$. Keeping $x_i(\tau)$ the same, this preserves feasibility of Eq. (6). Since the right-hand side of Eq. (5) does not decrease, the utility $u_i(\tau)$ does not decrease. This is also the mapping σ_i^a that maximizes $\Pr[\sigma_i^a]$.
2. If $\mu_i < m$, then $\mu_i(\sigma_i^b) < m$, which means $Q_i = 1$, and $p_i = \Pr[\sigma_i^a]$. If we use the mapping that maximizes $\Pr[\sigma_i^a]$ and keeping $x_i(\tau)$ the same, then this preserves the feasibility of Eq. (6) and does not decrease the RHS of Eq. (5), hence showing the utility $u_i(\tau)$ does not decrease.
3. If $\mu_i > \hat{m}$, then $\mu_i(\sigma_i^b) > \hat{m}$. In this case, $p_i = 1$, and $Q_i = \Pr[\sigma_i^a]$. Again, using the mapping that maximizes $\Pr[\sigma_i^a]$ and keeping $x_i(\tau)$ the same, this preserves the feasibility of Eq. (6) and does not decrease the RHS of Eq. (5), showing the utility $u_i(\tau)$ does not decrease.

Therefore, the resulting solution uses a maximal mapping and preserves Eqs. (5) and (6). \square

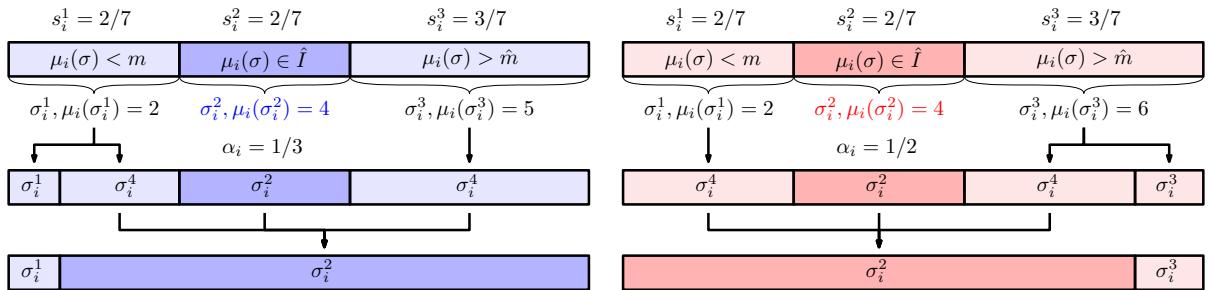


Figure 3: Illustration of the reduction process for two instances. In both instances, we have $s_i^1 = 2/7$, $s_i^2 = 2/7$, $s_i^3 = 3/7$; with $\mu_i(\sigma_i^1) = 2$ and $\mu_i(\sigma_i^2) = 4$. In the blue (left) instance, $\mu_i(\sigma_i^3) = 5$, so $\alpha_i = 1/3$. We absorb all the probability mass from σ_i^3 (i.e., $3/7$) together with $3/14$ probability of σ_i^1 into σ_i^2 . In the red (right) instance, $\mu_i(\sigma_i^3) = 6$ thus $\alpha_i = 1/2$. We absorb all the probability mass of σ_i^1 (i.e., $2/7$) with the same amount of probability mass in σ_i^3 into σ_i^2 .

4.3 Reduction to Network Flow and Approximate Majorization

For any agent i , the maximal mapping (that maximizes $\Pr[\sigma_i^a]$) is the solution to a linear program. Let y_v denote the probability that the value v maps to a signal in \hat{I} , and we have the following linear program.

$$\begin{aligned}
 & \text{maximize:} && \sum_v y_v \\
 & \text{subject to:} && \sum_v v \cdot y_v \geq m \cdot \sum_v y_v; \\
 & && \sum_v v \cdot y_v \leq \hat{m} \cdot \sum_v y_v; \\
 & && 0 \leq y_v \leq \Pr[v], \quad \forall v.
 \end{aligned} \tag{7}$$

Let the optimal objective value of Program (7) be β_i . If $\mu_i < m$, then we set $Q_i = 1$ and $p_i = \beta_i$; if $\mu_i > \hat{m}$, then we set $p_i = 1$ and $Q_i = \beta_i$; else we set $p_i = Q_i = 1$. Again, let $Q = \prod_i Q_i$.

We now rewrite the constraints and objective for the independent signaling scheme τ as follows. Note that $\{p_i\}_{i=1}^n$ and Q are calculated as above, and do not depend on the policy.

$$\begin{aligned}
m \cdot Q \cdot x_i(\tau) &= u_i(\tau), & \forall i \in [n]; \\
\sum_{i \in S} x_i(\tau) &\leq 1 - \prod_{i \in S} (1 - p_i(\tau)) & \forall S \subseteq [n].
\end{aligned} \tag{8}$$

Note that $\Omega_k^* = \sum_{\ell \geq 1} \rho_\ell \cdot \tau_\ell$, and after our transformation above, each τ_ℓ now faces the same constraints in [Program \(8\)](#). Therefore, we can set $x_i = \sum_{\ell \geq 1} \rho_\ell \cdot x_i(\tau_\ell)$, which will be feasible in [Program \(8\)](#). In order to reduce to network flow, we will relax the above program to the following:

$$\begin{aligned}
m \cdot Q \cdot x_i(\tau) &= u_i(\tau), & \forall i \in [n]; \\
\sum_{i=1}^n x_i(\tau) &\leq 1; \\
0 \leq x_i(\tau) &\leq p_i, & \forall i \in [n].
\end{aligned} \tag{9}$$

Recall from [Section 4.1](#) that we randomize over the projections $\{\Omega_k^*\}_{k=1}^K$ of the optimal signaling policy Ω^* , while preserving all utilities to a factor of K . For every $k \in \{1, 2, \dots, K\}$, consider the relaxed **SINGLE MEAN** policy constructed above. For notational convenience, we denote the p_i, Q_i in the corresponding maximal mapping as p_i^k, Q_i^k . Further denote the lower end-point of I_k as m^k , and the variable x_i by x_i^k . Putting together [Program \(9\)](#) for different values of k , it holds that the utilities $\{u_i\}_{i=1}^n$ of the policy Ω^* , after losing the factor of K , satisfy the following constraints.

$$\begin{aligned}
\frac{1}{K} \cdot \sum_{k=1}^K m^k \cdot Q^k \cdot x_i^k &= u_i, & \forall i \in [n]; \\
\sum_{i=1}^n x_i^k &\leq 1, & \forall k \in [K]; \\
0 \leq x_i^k &\leq p_i^k, & \forall i \in [n], k \in [K].
\end{aligned} \tag{10}$$

Network Flow and Majorization. Define $b^k := m^k \cdot Q^k / K$ and $y_i^k := b^k \cdot x_i^k$, and we have

$$\begin{aligned}
\sum_{k=1}^K y_k^i &= u_i, & \forall i \in [n]; \\
\sum_{i=1}^n y_i^k &\leq b^k, & \forall k \in [K]; \\
0 \leq y_i^k &\leq p_i^k \cdot b^k, & \forall i \in [n], k \in [K].
\end{aligned} \tag{11}$$

By the same argument as in [Section 3.2](#), the constraints above formulate a network flow problem. The agents are the sinks. Each interval \hat{I}_k is an intermediate node that is connected to the source with edges of capacity b^k . The node \hat{I}_k is connected to sink t_i with an edge of capacity $p_i^k \cdot b^k$. By [Theorem 3.3](#), this network flow problem has a 1-majorized solution (which is computable in polynomial time according to the discussion in [Section 3.1](#)). Consequently, there exist utilities $\{\hat{u}_i\}_{i=1}^n$ that are 1-majorized.

Final Algorithm. Our final algorithm will round the majorized solution analogous to [Section 3.2](#).

- For each $k \in \{1, 2, \dots, K\}$, consider the relaxed interval $\hat{I}_k = [m_k, \eta \cdot m_k)$, and compute the quantities $\{p_i^k, Q_i^k\}$ using [Program \(7\)](#).

- Compute the 1-majorized solution \hat{x} to [Program \(10\)](#).
- Pick a number $k \in \{1, 2, \dots, K\}$ uniformly at random.
- Given the realized values, compute the corresponding signals via the maximal mapping for \hat{I}_k and send the signals to the receiver. The receiver selects an agent as follows:
 - If any agent has posterior mean larger than the interval \hat{I}_k , select the agent with the highest mean and stop.
 - Consider the agents whose posterior mean lies in \hat{I}_k in a uniformly random order. For each agent i in this order, if it is reached, select it with probability \hat{x}_i^k / p_i^k and stop.
 - If no agent is selected so far, select the agent with the highest posterior mean.

Analysis. The same analysis as [Section 3.2](#) shows that for each agent i , the utility in the above algorithm is at least $\hat{u}_i/2$, where \hat{u}_i is the utility in the 1-majorized solution to [Program \(10\)](#). To see this, note that the probability of selecting the interval K is $1/K$. Conditioned on this event, the probability that no agent has posterior mean larger than \hat{I}_k is Q^k . Conditioned on these two events, the probability that the posterior mean of agent i both lies in \hat{I}_k and is selected is at least $\hat{x}_i^k/2$, and in this case, it yields utility at least m^k . The overall utility bound now follows by linearity of expectation.

Randomizing over SINGLE MEAN projections (see [Section 4.1](#)) loses a factor of $O(K) = O((\log V)/\varepsilon)$ in each utility, and hence our method gives a bicriteria $O((\log V)/\varepsilon)$ -majorized solution, completing the proof of [Theorem 2.8](#).

4.4 Extension to Asymmetric Utilities

So far, we have assumed that the receiver's selection criterion is based on the agents' signaled values, and the value of the chosen agent is the welfare the system generates. The work of [\[AK20, DTWZ24\]](#) considers a slightly different model, where though the agents signal their values and the receiver selects the agent whose mean posterior signaled value is largest, an individual agent's utility is 1 if selected and 0 otherwise. This corresponds to the utility perceived by the selfish agent, as opposed to the welfare generated by the system. We observe that the proofs for [Theorems 2.6](#) and [2.8](#) extends as is to this model, and sketch the changes needed.

To see that the proof of [Theorem 2.8](#) extends as is, we keep the construction of the mathematical program in [Section 4.2](#) the same, except we modify [Eq. \(5\)](#) to

$$u_i(\tau) = Q(\tau) \cdot x_i(\tau), \quad \forall i \in [n]. \quad (12)$$

which captures the utility being the same as the probability of being selected. For this modification, the proof of [Lemma 4.3](#) remains unchanged. Subsequently, [Program \(7\)](#) remains unchanged, while we change [Program \(10\)](#) to

$$\begin{aligned} \frac{1}{K} \cdot \sum_{k=1}^K Q^k \cdot x_i^k &= u_i, \quad \forall i \in [n]; \\ \sum_{i=1}^n x_i^k &\leq 1, \quad \forall k \in [K]; \\ 0 \leq x_i^k &\leq p_i^k, \quad \forall i \in [n], k \in [K]. \end{aligned} \quad (13)$$

This again reduces to network flow by defining $b^k := Q^k/K$ and $y_i^k := b^k \cdot x_i^k$, yielding exactly [Program \(11\)](#). This completes the proof of [Theorem 2.8](#) for this setting.

Analogously, the proof of [Theorem 2.6](#) extends to this setting by changing the relaxation to replace the value v_j by 1 in the first (utility) constraint [Eq. \(1\)](#). It is easy to check that the resulting program this also reduces to network flow. The rest of the analysis goes through as is, except for replacing v_j by 1 in [Eq. \(4\)](#). This shows a 2-majorized solution.

5 Lower Bound on Majorization: Proof of [Theorem 2.9](#)

We first assume that the receiver is maximizing the utilitarian welfare exactly, and it will become clear that the same lower bound holds if the receiver selects an *arbitrary* function of the posterior means – in particular, if it is a $(1 + \varepsilon)$ -approximate utilitarian welfare maximizer.

Let there be n agents. Each agent i has value $i + 1$ with probability $1/i$ and has value 1 with probability $1 - 1/i$. Note that the expected value of agent i is $\mu_i = 2$, the same for each i . For any $k \in [n]$, denote $\sum_{i=k}^n 1/(i + 1)$ by H_{-k} . We first give a lower bound on the sum of the smallest k agents' utilities.

Lemma 5.1. *For any $k \in [n]$, there exists a signaling policy S_k , such that if the utilities of the agents are arranged in ascending order, the prefix of the first k utilities sums to*

$$\frac{1}{2} \cdot \frac{k \cdot (k + 3)}{k + (k + 1) \cdot H_{-(k+1)}}.$$

In order to prove this lemma, we construct the signaling policy S_k as follows. For each agent i with $i \in \{k + 1, \dots, n\}$, if $v_i = i + 1$, then with probability x_i , the sender sends a signal s_i to the receiver, and otherwise the sender sends a signal \bar{s}_i for this agent. For $i \leq k$, the sender reveals the true value of agent i . Note that this mapping rule is independent for each agent.

Note that conditioned on receiving signal s_i , we have $\mu_i(s_i) = i + 1$. On the other hand, conditioned on receiving \bar{s}_i , we have

$$\mu_i(\bar{s}_i) = \frac{\left(1 - \frac{1}{i}\right) \cdot 1 + \frac{1}{i} \cdot (1 - x_i) \cdot (i + 1)}{1 - \frac{1}{i} \cdot x_i} = \frac{2 - x_i - \frac{x_i}{i}}{1 - \frac{x_i}{i}} \leq 2.$$

Therefore, if at least one signal s_i is received, the receiver will select the agent i , $i \in [k + 1, n]$ with the largest index. On the other hand, if no signal s_i is received, the receiver selects the unique agent $i \in [k]$ with the highest revealed value.

For each $i \in \{k + 1, \dots, n\}$, we set

$$x_i = \frac{\frac{i}{i+1}}{\frac{k}{k+1} + H_{-(k+1)} - H_{-(i+1)}}.$$

The intuition behind this formula is to have $U_i = U_{i+1}$ for any $k \leq i \leq n - 1$. We obtain the utility of all agents as follows.

Lemma 5.2. *It holds that*

$$U_j(S_k) = \begin{cases} \frac{j+1}{k+1} \cdot \frac{1}{\frac{k}{k+1} + H_{-(k+1)}} & j \leq k; \\ \frac{1}{\frac{k}{k+1} + H_{-(k+1)}} & j \geq k + 1. \end{cases}$$

Proof. By the construction of S_k , for any $k + 1 \leq j \leq n$, the receiver will receive signal s_j and no signal s_i for $i > j$ with probability

$$x_j \cdot \frac{1}{j} \cdot \prod_{i=j+1}^n \left(1 - \frac{x_i}{i}\right).$$

Since agent j receives utility $j + 1$ if and only if signal s_j is sent and no signal s_i for $i > j$ is sent, we have

$$\begin{aligned}
U_j(S_k) &= (j+1) \cdot x_j \cdot \frac{1}{j} \cdot \prod_{i=j+1}^n \left(1 - \frac{x_i}{i}\right) = \frac{j+1}{j} \cdot x_j \cdot \prod_{i=j+1}^n \left(1 - \frac{\frac{1}{i+1}}{\frac{k}{k+1} + H_{-(k+1)} - H_{-(i+1)}}\right) \\
&= \frac{1}{\frac{k}{k+1} + H_{-(k+1)} - H_{-(j+1)}} \cdot \prod_{i=j+1}^n \left(\frac{\frac{k}{k+1} + H_{-(k+1)} - H_{-(i+1)} - \frac{1}{i+1}}{\frac{k}{k+1} + H_{-(k+1)} - H_{-(i+1)}}\right) \\
&= \frac{1}{\frac{k}{k+1} + H_{-(k+1)} - H_{-(j+1)}} \cdot \prod_{i=j+1}^n \left(\frac{\frac{k}{k+1} + H_{-(k+1)} - H_{-i}}{\frac{k}{k+1} + H_{-(k+1)} - H_{-(i+1)}}\right) = \frac{1}{\frac{k}{k+1} + H_{-(k+1)}}.
\end{aligned}$$

Next we consider agent j for $j \leq k$. By a similar telescoping as above, the probability that no signal s_i with $i > k$ is sent (denoted by q_R) is

$$q_R = \prod_{i=k+1}^n \left(1 - \frac{\frac{1}{i+1}}{\frac{k}{k+1} + H_{-(k+1)} - H_{-(i+1)}}\right) = \frac{\frac{k}{k+1} + H_{-(k+1)} - H_{-(k+1)}}{\frac{k}{k+1} + H_{-(k+1)}} = \frac{\frac{k}{k+1}}{\frac{k}{k+1} + H_{-(k+1)}}.$$

For any agent j with $j \leq k$, the probability that it has the largest revealed value among agents $\{1, 2, \dots, k\}$ is

$$\frac{1}{j} \cdot \prod_{i=j+1}^k \left(1 - \frac{1}{i}\right) = \frac{1}{k}.$$

Therefore, the expected utility of agent j is

$$U_j(S_k) = \frac{j+1}{k} \cdot q_R = \frac{j+1}{k} \cdot \frac{\frac{k}{k+1}}{\frac{k}{k+1} + H_{-(k+1)}} = \frac{j+1}{k+1} \cdot \frac{1}{\frac{k}{k+1} + H_{-(k+1)}}. \quad \square$$

Completing the proof of Lemma 5.1. We apply the signaling scheme S_k defined above. By Lemma 5.2, we have $U_i(S_k) \leq U_j(S_k)$ for all $1 \leq i \leq j \leq n$. Thus the k agents with smallest expected utilities are agents $1, 2, \dots, k$. Their utilities sum up to

$$\sum_{i=1}^k U_i(S_k) = \frac{1}{2} \cdot \frac{k \cdot (k+3)}{k + (k+1) \cdot H_{-(k+1)}},$$

completing the proof. \square

Proof of Theorem 2.9. Since $k \geq 1$, we have

$$\frac{1}{2} \cdot \frac{k \cdot (k+3)}{k + (k+1) \cdot H_{-(k+1)}} > \frac{1}{2} \cdot \frac{k \cdot (k+3)}{(k+1) + (k+1) \cdot H_{-(k+1)}} \geq \frac{1}{2} \cdot \frac{k+1}{1 + H_{-(k+1)}} =: R_k.$$

The quantity R_k is a lower bound on the sum of the k smallest utilities in the signaling policy S_k . Fix any signaling policy S' . Suppose that each agent i is selected with probability w_i in S' . If S' is α -majorized, the sum of the utilities of agents in the set $\{1, 2, \dots, k\}$ must be at least R_k/α . Therefore, we have the following set of inequalities

$$\sum_{i=1}^k (i+1) \cdot w_i \geq \frac{1}{\alpha} \cdot R_k, \quad \forall k \in \{1, 2, \dots, n\}.$$

We now want to lower bound $\sum_{i=1}^n w_i$. The optimal way to assign the probabilities $\{w_i\}_{i=1}^n$ is to make all the inequalities satisfied with equality. Therefore, we have

$$\begin{aligned}
\sum_{i=1}^n w_i &\geq \frac{1}{2\alpha} \cdot R_1 + \frac{1}{\alpha} \cdot \sum_{i=2}^n \left(\frac{R_i - R_{i-1}}{i+1} \right) \\
&\geq \frac{1}{\alpha} \cdot \sum_{i=1}^n \frac{R_i}{(i+1) \cdot (i+2)} = \frac{1}{2\alpha} \cdot \sum_{i=1}^n \frac{1}{(i+2) \cdot (1 + H_{-(i+1)})} \\
&\geq \frac{1}{2\alpha} \cdot \sum_{i=1}^n \frac{1}{(i+2) \cdot (1 + \log(n+1) - \log(i+1))} \quad (\text{Since } H_{-(i+1)} \leq \log\left(\frac{n+1}{i+1}\right)) \\
&\geq \frac{1}{2\alpha} \cdot \sum_{i=1}^n \frac{1}{(\frac{3}{2}i + \frac{3}{2}) \cdot (1 + \log(n+1) - \log(i+1))} \\
&\geq \frac{1}{3\alpha} \cdot \int_1^{n+1} \frac{1}{x \cdot (1 + \log(n+1) - \log x)} dx \\
&\quad (\text{Since } \frac{1}{x \cdot (1 + \log(n+1) - \log x)} \text{ is monotonically decreasing on } [1, n+1]) \\
&= \frac{1}{3\alpha} \cdot \left[-\log(1 + \log(1 + n) - \log x) \Big|_{x=1}^{n+1} \right] = \frac{1}{3\alpha} \cdot \log(1 + \log(1 + n)).
\end{aligned}$$

Combining the inequality above with $\sum_{i=1}^n w_i \leq 1$, we have

$$\alpha \geq \frac{\log(1 + \log(1 + n))}{3} > \frac{1}{3} \cdot \log \log(n+1) = \frac{1}{3} \cdot \log \log V. \quad \square$$

Since the second part of our argument only requires the fact that at most one agent is selected and is independent of the receiver's selection rule, this lower bound also applies if the receiver can arbitrarily select the agent, or (in particular) if the receiver is an approximate utilitarian welfare maximizer.

6 Open Directions

Our work points to several interesting future directions. First, we assume independent (or decentralized) mapping of values to signals, motivated hiring and selection applications. Can a similar existence result as in [Section 4](#) extend to the setting where the sender can correlate signals from different agents? We note however that the two settings are somewhat incomparable, and it may very well be that the correlated setting is computationally simpler for a given fairness function [\[DX16\]](#), while the independent case is easier from an approximate majorization perspective.

Secondly, our lower bound in [Section 5](#) holds for distributions with large variance. Is there an $O(1)$ -majorized policy under a more benign assumption on distributions, such as the monotone hazard rate (MHR) assumption? Next, we assumed a single agent is finally selected. What if the receiver selects the top k agents according to the posterior mean? In this case, it is open if there a polynomial-time algorithm for a given fairness function, in both the cases where the sender can correlate signals and when it sends independent signals. Further, it is an open question to extend our majorization result even to the case when $k = 2$.

Finally, we assume agents are not strategic in revealing information, and our results can be viewed as the limits of fairness that is achievable even if agents follow a prescribed policy. It is an interesting direction to study equilibria and price of anarchy when agents reveal information strategically, building on [\[AK20, DTWZ24\]](#).

References

- [AK20] Pak Hung Au and Keiichi Kawai. Competitive information disclosure by multiple senders. *Games and Economic Behavior*, 119:56–78, 2020.
- [BBM15] Dirk Bergemann, Benjamin Brooks, and Stephen Morris. The limits of price discrimination. *American Economic Review*, 105(3):921–957, 2015.
- [BM04] Marianne Bertrand and Sendhil Mullainathan. Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. *American Economic Review*, 94(4):991–1013, 2004.
- [BM19] Dirk Bergemann and Stephen Morris. Information design: A unified perspective. *Journal of Economic Literature*, 57(1):44–95, 2019.
- [BMSW24] Siddhartha Banerjee, Kamesh Munagala, Yiheng Shen, and Kangning Wang. Fair price discrimination. In *Proceedings of the 2024 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2679–2703. SIAM, 2024.
- [BTXZ21] Yakov Babichenko, Inbal Talgam-Cohen, Haifeng Xu, and Konstantin Zabarnyi. Regret-minimizing bayesian persuasion. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*, page 128. ACM, 2021.
- [CDHW20] Rachel Cummings, Nikhil R. Devanur, Zhiyi Huang, and Xiangning Wang. Algorithmic price discrimination. In *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2432–2451. SIAM, 2020.
- [CDW12] Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. Optimal multi-dimensional mechanism design: Reducing revenue to welfare maximization. In *Proceedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 130–139. IEEE Computer Society, 2012.
- [CH14] Archishman Chakraborty and Rick Harbaugh. Persuasive puffery. *Marketing Science*, 33(3):382–400, 2014.
- [CMV20] L. Elisa Celis, Anay Mehrotra, and Nisheeth K. Vishnoi. Interventions for ranking in the presence of implicit bias. In *Proceedings of the 2020 ACM Conference on Fairness, Accountability, and Transparency (FAT*)*, pages 369–380. ACM, 2020.
- [CS19] Deeparnab Chakrabarty and Chaitanya Swamy. Approximation algorithms for minimum norm and ordered optimization problems. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 126–137. ACM, 2019.
- [DGV08] Brian C. Dean, Michel X. Goemans, and Jan Vondrák. Approximating the stochastic knapsack problem: The benefit of adaptivity. *Mathematics of Operations Research*, 33(4):945–964, 2008.
- [DKKS24] Siddartha Devic, Aleksandra Korolova, David Kempe, and Vatsal Sharan. Stability and multi-group fairness in ranking with uncertain predictions. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*. PMLR, 2024.
- [DKQ16] Shaddin Dughmi, David Kempe, and Ruixin Qiang. Persuasion with limited communication. In *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*, pages 663–680. ACM, 2016.

[DR89] Bhaskar Dutta and Debraj Ray. A concept of egalitarianism under participation constraints. *Econometrica*, 57(3):615–635, 1989.

[DREM13] Ezekiel J. Dixon-Román, Howard T. Everson, and John J. McArdle. Race, poverty and SAT scores: Modeling the influences of family income on black and white high school students’ SAT performance. *Teachers College Record*, 115(4):1–33, 2013.

[DTWZ24] Zhicheng Du, Wei Tang, Zihe Wang, and Shuo Zhang. Competitive information design with asymmetric senders. In *Proceedings of the 25th ACM Conference on Economics and Computation (EC)*, 2024.

[Dug17] Shaddin Dughmi. Algorithmic information structure design: a survey. *SIGecom Exchanges*, 15(2):2–24, 2017.

[DX16] Shaddin Dughmi and Haifeng Xu. Algorithmic bayesian persuasion. In Daniel Wichs and Yishay Mansour, editors, *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 412–425. ACM, 2016.

[EGGL20] Vitalii Emelianov, Nicolas Gast, Krishna P. Gummadi, and Patrick Loiseau. On fair selection in the presence of implicit variance. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*, pages 649–675. ACM, 2020.

[GB21] David García-Soriano and Francesco Bonchi. Maxmin-fair ranking: Individual fairness under group-fairness constraints. In *Proceedings of 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, pages 436–446. ACM, 2021.

[GM06] Ashish Goel and Adam Meyerson. Simultaneous optimization via approximate majorization for concave profits or convex costs. *Algorithmica*, 44(4):301–323, 2006.

[GMP05] Ashish Goel, Adam Meyerson, and Serge A. Plotkin. Approximate majorization and fair online load balancing. *ACM Transactions on Algorithms*, 1(2):338–349, 2005.

[GMS10] Sudipto Guha, Kamesh Munagala, and Peng Shi. Approximation algorithms for restless bandit problems. *Journal of the ACM*, 58(1):3:1–3:50, 2010.

[GR00] Claudia Goldin and Cecilia Rouse. Orchestrating impartiality: The impact of “blind” auditions on female musicians. *American Economic Review*, 90(4):715–741, 2000.

[HLP34] Godfrey Harold Hardy, John Edensor Littlewood, and George Pólya. *Inequalities*. Cambridge university press, 1934.

[Kar32] Jovan Karamata. Sur une inégalité relative aux fonctions convexes. *Publications de l’Institut Mathématique*, 1(1):145–147, 1932.

[KG11] Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.

[KK06] Amit Kumar and Jon M. Kleinberg. Fairness measures for resource allocation. *SIAM Journal on Computing*, 36(3):657–680, 2006.

[KR18] Jon M. Kleinberg and Manish Raghavan. Selection problems in the presence of implicit bias. In *Proceedings of the 9th Conference on Innovations in Theoretical Computer Science Conference (ITCS)*, pages 33:1–33:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018.

- [MC21] Anay Mehrotra and L. Elisa Celis. Mitigating bias in set selection with noisy protected attributes. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT)*, pages 237–248. ACM, 2021.
- [Meg74] Nimrod Megiddo. Optimal flows in networks with multiple sources and sinks. *Mathematical Programming*, 7:97–107, 1974.
- [SKJ21] Ashudeep Singh, David Kempe, and Thorsten Joachims. Fairness in ranking under uncertainty. In *Proceedings of the 35th Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 11896–11908, 2021.
- [SWZ⁺23] Zeyu Shen, Zhiyi Wang, Xingyu Zhu, Brandon Fain, and Kamesh Munagala. Fairness in the assignment problem with uncertain priorities. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 188–196. ACM, 2023.
- [Tam95] Arie Tamir. Least majorized elements and generalized polymatroids. *Mathematics of Operations Research*, 20(3):583–589, 1995.
- [Vei71] Arthur F. Veinott Jr. Least d -majorized network flows with inventory and statistical applications. *Management Science*, 17(9):547–567, 1971.
- [XRDT15] Haifeng Xu, Zinovi Rabinovich, Shaddin Dughmi, and Milind Tambe. Exploring information asymmetry in two-stage security games. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, pages 1057–1063. AAAI Press, 2015.