

Quantile Agent Utility and Implications to Randomized Social Choice

Ioannis Caragiannis and Sanjukta Roy

Abstract

We initiate a novel direction in randomized social choice by proposing a new definition of agent utility for randomized outcomes. Each agent has a preference over all outcomes and a *quantile* parameter. Given a *lottery* over the outcomes, an agent gets utility from a particular *representative*, defined as the least preferred outcome that can be realized so that the probability that any worse-ranked outcome can be realized is at most the agent’s quantile value.

In contrast to other utility models that have been considered in randomized social choice (e.g., stochastic dominance, expected utility), our *quantile agent utility* compares two lotteries for an agent by just comparing the representatives, as is done for deterministic outcomes.

We revisit questions in randomized social choice using the new utility definition. We study the compatibility of efficiency and strategyproofness for randomized voting rules, efficiency and fairness for randomized one-sided matching mechanisms, and efficiency, stability, and strategyproofness for lotteries over two-sided matchings. In contrast to well-known impossibilities in randomized social choice, we show that satisfying the above properties simultaneously is possible.

1 Introduction

Randomized mechanisms are pivotal in social choice and matching theory, offering innovative solutions to challenges such as fairness, efficiency, and strategyproofness. Recent global events have further underscored their importance. For instance, during the COVID-19 pandemic, randomized algorithms played a critical role in vaccine allocation to ensure equitable distribution across diverse populations. Today, randomized mechanisms are increasingly employed to assign limited resources such as hospital beds, educational opportunities, and public housing to applicants in a manner that balances fairness and efficiency. In voting theory, randomized mechanisms have been used to break ties or select representatives fairly in multi-agent decision-making environments, such as participatory budgeting or citizen assemblies, where no deterministic method can satisfy all stakeholders.

Traditional approaches in *randomized social choice* have often grappled with inherent incompatibilities. For instance, achieving important axioms such as both efficiency and strategyproofness simultaneously has been shown to be fundamentally problematic [9]. Gibbard’s theorem identifies (randomized versions of) dictatorships as the only voting systems that satisfy the two axioms simultaneously [24]. However, such a negative result strongly depends on assumptions regarding how the voters (or agents) evaluate the voting outcome.

In this paper, we propose a novel definition of agent utility that is specifically tailored to randomized outcomes. Our framework shifts the focus from traditional comparisons, such as stochastic dominance or expected utility (e.g., see [14]), to defining utility based on an agent’s least preferred outcome that meets their probabilistic expectations. By adopting this perspective, we aim to reconcile the often conflicting objectives of efficiency, stability, and strategyproofness in social choice mechanisms.

More specifically, the study of axiomatic properties in social choice theory has been traditionally based on how agents compare different outcomes. For example, in a classical voting scenario, one would like to aggregate preferences over alternatives submitted by the voters (or agents) into a single winning alternative. An agent’s preference (usually, a strict ranking) over the alternatives determines how the

agent evaluates different deterministic outcomes that may be returned by a voting rule and compares them to each other. Similarly, in a two-sided matching instance, we are given two disjoint sets of agents (e.g., men and women in a stable marriage instance), with each agent having a preference on the agents of the other side. Important properties like stability of matchings are defined using the agents' preferences.

The relatively recent trend of randomized social choice aims to study the role of randomization in the procedures above. How do agents evaluate the outcome of randomized voting rules, random assignments in one-sided matching instances, or fractional two-sided matchings? For example, in the voting scenario mentioned above, a randomized voting rule takes as input a profile of agent preferences and returns a random alternative according to a probability distribution. In other words, a randomized voting rule takes as input a profile of agent preferences (rankings over the alternatives) and returns a *lottery* over the alternatives.

The important question now is: how can an agent decide which among two lotteries is better using her preference ranking of the alternatives? The answer is not trivial anymore. For example, how can an agent with preference $a \succ b \succ c$ over three alternatives a , b , and c compare the lottery that returns equiprobably the three alternatives with the lottery that returns alternative b with probability 1? There have been several answers in the literature. One approach is to assume that the agent has a valuation function that assigns a cardinal utility for each alternative that respect the preference relation of the agent for the alternatives. In our example, the utility for alternative a has to be higher than that for alternative b , which in turn has to be higher than that for alternative c . The two lotteries can now be compared according to the expected utility of the agent for the alternative they return. This approach requires the existence of underlying cardinal utilities, which the agents should be able to compute; this is not always realistic.

Another prominent approach in randomized social choice is to compare lotteries using the notion of stochastic dominance. According to it, a lottery is better than another for an agent if it yields higher expected utility for her for every possible underlying utilities for the alternatives that respect her preference ranking. This is a very stringent definition which may result in two lotteries being incomparable. For example, the comparison between the two lotteries mentioned above gives a different outcome for the underlying utilities $(1, 0.2, 0)$ and $(1, 0.8, 0)$ for alternatives a , b , and c , respectively. The expected utilities for the two lotteries are 0.4 and 0.2 for the former and 0.6 and 0.8 for the latter.

In this work, we present a new agent utility model for the comparison of lotteries. Each agent has a *quantile parameter* (a scalar between 0 and 1), and associates each lottery over alternatives with the highest-ranked alternative that has the following properties: it has positive probability and it is such that the less preferred alternatives have total probability that does not exceed the quantile parameter. We use the term *representative* to refer to this alternative. Then, the comparison of two lotteries by an agent is simply the comparison of their representative alternatives according to the agent's preference. We use the term *quantile agent utility* to refer to our new utility model.

Our quantile agent utility model has several advantages. First, the comparison between lotteries can take place in the very same way as in the case of deterministic outcomes. Second, as we will see, it turns out that stochastic dominance is as stringent in the comparison between two lotteries as our utility definition would be by considering all possible quantile parameters. And, of course, there are no underlying cardinal utilities for the alternatives that the agent needs to be able to evaluate the lotteries.

The different quantile parameters can model agents with different risk levels in their interpretation of a lottery. As should be expected, different quantile parameter values can lead to different results. For example, if all agents have their quantile parameter equal to 0, we usually recover the impossibilities for deterministic mechanisms. If, instead, all quantile parameters are close to 1, trivial mechanisms like uniform lotteries are usually ideal. Intermediate quantile parameter values, possibly different for each agent, allow for a spectrum of interesting results.

1.1 Our contribution

The definition of the quantile agent utility model is our main conceptual contribution (see Section 2). More importantly, we explore its implications to three social choice settings: voting, one-, and two-sided matchings.

Voting. In Section 3, we consider the classical ranking-based voting scenario, in which a voting rule aims to select a winning alternative (candidate) taking as input rankings of the available alternatives that are submitted as votes by a set of agents (voters). The voting rule is randomized, and the agents have quantile utilities. We adapt the definition of (Pareto) efficiency in this context and first observe that no lottery can be efficient for all vectors of quantile parameters for the agents unless all agents agree on their top-ranked alternative. This justifies our choice to focus on scenarios in which each agent has their quantile parameter, which in turn can be used in the definition of the lottery returned by the voting rule. We then adapt the notion of strategyproofness and explore its relation to monotonicity in the new context for the case of two alternatives. We conclude this section with the more intricate case of three alternatives and profiles in which all agents have the same quantile parameter. For high values of the quantile parameter, we present non-dictatorial voting rules that are efficient and strategyproof. For small values, we show that deterministic dictatorships are the only efficient and strategyproof rules. Extending our characterization for all possible quantile parameter vectors is left as an (apparently challenging) important open problem.

One-sided matching. In Section 4, we consider scenarios with a set of agents that have to be matched (in a one-to-one manner) to a set of items of equal size. Each agent has a ranking of the items representing her preference over them and a quantile parameter. We consider randomized matching mechanisms that return lotteries over matchings. These lotteries define a representative item for each agent; the position of this item in the agent’s ranking indicates how desirable the lottery is for the agent. We begin by considering a simple variant of social welfare in this context and show that optimal lotteries over matchings can be easily computed. We also discuss more elaborate definitions of social welfare, which seem to be more challenging to optimize. Again, we adapt the notions of efficiency and strategyproofness and demonstrate how to adapt serial dictatorship in our context using linear programming. Next, motivated by the fair division literature, we consider fairness properties. We define the analog of proportionality and show how to adapt serial dictatorships to get a matching mechanism that computes efficient and proportional lotteries. Unfortunately, our rule is not strategyproof. Whether matching mechanisms that are simultaneously efficient, strategyproof, and proportional exist is another open problem for the one-sided matching setting. Furthermore, we give a natural definition of envy-freeness. Our definitions guarantee that envy-freeness implies proportionality (in sync with the classical fair division literature). We observe that envy-freeness may not be consistent with efficiency; whether envy-free and efficient lotteries can be computed in polynomial time is another open problem for one-sided matchings.

Two-sided matching. We also consider another matching scenario with two disjoint sets of agents in which each agent has a quantile parameter, and every agent in one set has a ranking of the agents in the other set as their preference. We briefly consider a variant of social welfare but move quickly to redefining the notion of stability. We present structural observations for stable lotteries over matchings and discuss their relations with integral stable matchings and ordinal stability from the classical stable matchings literature. We demonstrate that stability and efficiency are not compatible for small values of the quantile parameters and present possibility results for a restricted version of efficiency, strategyproofness, and stability when the quantile parameter is above $1/2$. Determining the range of quantile parameters that allow for the (unrestricted version of the) three properties simultaneously is the main open problem for two-sided matchings under agents with quantile utilities. Due to space constraints, these results are presented in Appendix C.

1.2 Related work

Lotteries over alternatives were first formally studied by Zeckhauser [45], Fishburn [20], and Intriligator [28]. Randomization in voting processes have been considered to address impossibility results in deterministic mechanisms, such as Arrow’s Impossibility Theorem [6]. One of the most famous results in this area is by Gibbard [23] and Satterthwaite [40], who showed that deterministic strategyproof mechanisms must be dictatorial, inspiring the use of randomization to achieve strategyproofness and fairness. Duggan and Schwartz [19] generalized it for lotteries. Bogomolnaia and Moulin [13] proposed the Random Priority (RP) mechanism, ensuring fairness and strategyproofness in resource allocation and voting. Procaccia [36] designed randomized voting rules that are approximations of score-based deterministic voting rules to achieve strategyproofness. Sato [39] showed non-existence of desirable strategyproof deterministic rules when agent preferences are extended to be weak orders over subsets of alternatives, either by considering the best or the worst alternative in the subsets. These two extensions to weak order are analogous to agents having quantile values either close to one or zero, respectively. Aziz et al. [9] showed that the efficiency and strategyproofness of a system can vary depending on how preferences over alternatives are extended to preferences over lotteries. They considered preference extensions using stochastic dominance (SD), pairwise comparisons (PC), bilinear dominance (BD), and sure-thing principle (ST). They showed that while random serial dictatorships are PC-strategyproof, they only achieve ex post efficiency, strict maximal lotteries are both PC-efficient and ST-strategyproof, and show multiple incompatibility results. See also the survey of related results by Brandt [14].

In resource allocation problems consisting of agents that need to be matched to items, envy-freeness [21, 44] and proportionality [41] have been widely studied in both economics and computer science as measures of fairness. Hylland and Zeckhauser [27] focused on the probabilistic allocations of indivisible items without relying on a priority structure and showed that this approach is more efficient than randomization over deterministic methods based on priority. Abdulkadiroğlu and Sönmez [1] studied lotteries in the house allocation problem, which is another variation of the problem of matching under one-sided preferences. Aziz [7] explored the advantages of applying randomization in social choice scenarios, such as fair division, and examined the associated challenges. Caragiannis et al. [16] studied interim envy-freeness for lotteries. For resource allocation, obtaining both ex-post and ex-ante guarantees simultaneously has also been considered [11, 10].

More discussion on earlier work related to two-sided and stable matchings is presented in Appendix A.

2 The quantile agent utility model

We define the quantile utility model as follows. Let O denote a finite set of options. Consider an agent equipped with a strict preference order $\succ$ among the options of O and a quantile parameter $h \in [0, 1)$. Given a *lottery* (or probability distribution) x over the options in O , we call option $o \in O$ the h -quantile representative¹ of the agent if $\sum_{o \succ o'} x(o') \leq h < \sum_{o \succeq o'} x(o')$. The h -quantile representative of the agent in lottery x is denoted by $\text{rep}(x, \succ, h)$. For example, if an agent has quantile parameter $h = 0$, then the representative is the lowest ranked alternative that has positive probability; for $h = 1/2$, the representative is the median.

Given two lotteries x and y over the options in O , we say that the agent prefers x to y if and only if $\text{rep}(x, \succ, h) \succ \text{rep}(y, \succ, h)$. We say that the agent weakly prefers x to y if and only if $\text{rep}(x, \succ, h) \succeq \text{rep}(y, \succ, h)$ (i.e., either $\text{rep}(x, \succ, h) \succ \text{rep}(y, \succ, h)$ or $\text{rep}(x, \succ, h) = \text{rep}(y, \succ, h)$).

We now explore the connection of our lottery extensions to the notion of *stochastic dominance*; we use the definition of [14].

¹We can define the 1-quantile representative of an agent to be the highest-ranked option that has positive probability to be returned by the lottery. However, to keep the exposition simple and the proofs easy to follow, we have decided to restrict quantile parameters in $[0, 1)$.

Definition 1 (stochastic dominance). We say that lottery x over the set of options O is not stochastically dominated by lottery y with respect to a preference $\succ$ over the options (and write $x \succsim^{sd} y$) if and only if

$$\sum_{o': o \succeq o'} x(o') \leq \sum_{o': o \succeq o'} y(o')$$

for every option $o \in O$.

It turns out that stochastic dominance is much more stringent than our new notion.

Theorem 2. Let O be a set of options, x and y be lotteries over the options in O , and $\succ$ be a preference order over O . Then, $x \succsim^{sd} y$ if and only if $\text{rep}(x, \succ, h) \succeq \text{rep}(y, \succ, h)$ for every $h \in [0, 1)$.

Proof. Assume that $x \succsim^{sd} y$ and, for the sake of contradiction, $a = \text{rep}(y, \succ, h^*) \succ \text{rep}(x, \succ, h^*) = b$ for some $h^* \in [0, 1)$. By the definition of the h^* -quantile representative $\text{rep}(x, \succ, h^*)$, we get

$$\sum_{o': b \succeq o'} x(o') > h^*. \quad (1)$$

Furthermore,

$$h^* \geq \sum_{o': a \succ o'} y(o') \geq \sum_{o': b \succeq o'} y(o'). \quad (2)$$

In this last equation, the first inequality follows from the definition of the h^* -quantile representative and the second one follows from the fact that $a \succ b$. By equations (1) and (2), we get $\sum_{o': b \succeq o'} x(o') > \sum_{o': b \succeq o'} y(o')$, contradicting the assumption $x \succsim^{sd} y$. This completes the proof of the “only if” part of the statement in Theorem 2.

Now assume that $\text{rep}(x, \succ, h) \succeq \text{rep}(y, \succ, h)$ for every $h \in [0, 1)$. For the sake of contradiction, assume that $x \not\succsim^{sd} y$. Then, there exists an option $o \in O$ such that

$$\sum_{o': o \succeq o'} x(o') > \sum_{o': o \succeq o'} y(o') \quad (3)$$

and

$$\sum_{o': o \succ o'} x(o') \leq \sum_{o': o \succ o'} y(o'). \quad (4)$$

Define

$$h^* = \sum_{o': o \succeq o'} y(o') \quad (5)$$

and notice that $h^* < 1$, due to equation (3). Equations (3) and (4) imply that $x(o) > 0$ and equation (5) yields $\text{rep}(x, \succ, h^*) = o$. Now, let a be the least preferred option in O that satisfies $a \succ o$ and $y(a) > 0$. Such an option exists since inequality (3) implies that $\sum_{o': o \succ o'} y(o') < 1$. We obtain that $\text{rep}(y, \succ, h^*) = a \succ o = \text{rep}(x, \succ, h^*)$, a contradiction. The “if” part of Theorem 2 follows, completing the proof. $\square$

3 Voting

We first study the quantile agent utility assumption in a voting setting with a set A of m alternatives and a set N of n agents (voters). Each agent $i \in N$ has a quantile parameter $h_i \in [0, 1)$. We use $\mathbf{h} = (h_1, \dots, h_n)$ to denote the vector of the quantile parameters. A voting profile $\succ = (\succ_1, \dots, \succ_n)$ consists of the preference $\succ_i$ of each agent; $\succ_i$ is a strict ordering over the alternatives. We use the term *voting instance* and identify the above scenario with the tuple $(N, A, \succ, \mathbf{h})$. We adapt the notion of efficiency for lotteries and agents with quantile utilities as follows.

Definition 3 (efficiency). *Given a voting instance $(N, A, \succ, \mathbf{h})$, a lottery x over the alternatives of A is efficient if there is no other lottery y such that $\text{rep}(y, \succ_i, h_i) \succeq_i \text{rep}(x, \succ_i, h_i)$ for every agent $i \in N$ and there is an agent $i^* \in N$ such that $\text{rep}(y, \succ_{i^*}, h_{i^*}) \succ_{i^*} \text{rep}(x, \succ_{i^*}, h_{i^*})$.*

Our next result indicates that efficient lotteries are well-defined only for specific quantile parameter vectors.

Theorem 4. *Consider a set N of n agents, a set A of alternatives, and a preference profile $\succ$ with the preferences of the agents in N over the alternatives in A . There is a lottery that is efficient for the voting instance $(N, A, \succ, \mathbf{h})$ for every quantile parameter vector $\mathbf{h} \in [0, 1]^n$ if and only if all agents have the same alternative ranked first.*

Proof. To prove the “if” part of the theorem, assume that the same alternative $a \in A$ is in the top position of the preference $\succ_i$ of every agent $i \in N$. Then, the lottery x with $x(a) = 1$ and $x(o) = 0$ for $o \in A \setminus \{a\}$ makes alternative a the representative of each agent i for every value of her quantile parameter $h_i \in [0, 1]$. Clearly, the lottery x is efficient.

For the “only if” part, let L be the set of top-ranked alternatives and assume that $|L| \geq 2$. First, consider the quantile parameter vector $\mathbf{h}^1 = (0, 0, \dots, 0)$. We claim that a lottery is efficient for $(N, A, \succ, \mathbf{h}^1)$ only if it returns one of the alternatives with probability 1. Indeed, assume otherwise and let $S \subseteq A$ be the set of alternatives with positive probability under lottery x , i.e., $x(o) > 0$ for $o \in S$ and $x(o) = 0$ for $o \in A \setminus S$, with $|S| \geq 2$. By definition, the representative $\text{rep}(x, \succ_i, 0)$ of agent $i \in N$ is the alternative in S that is ranked lowest among the alternatives in S . Then, lottery x is dominated by the lottery x' defined as $x'(a) = 1$ for the alternative $a \in S$ that is ranked higher than any other alternative in S by agent 1 and $x'(o) = 0$ for $o \in A \setminus \{a\}$. In this way, agent 1 has a strictly better representative under x' compared to x , while the representative under x' of any other agent is at least as good as the one under x .

Next, set $\ell = |L| \geq 2$ and consider the quantile parameter vector $\mathbf{h}^2 = (1 - 1/\ell, 1 - 1/\ell, \dots, 1 - 1/\ell)$. The lottery y with $y(o) = 1/\ell$ for $o \in L$ and $y(o) = 0$ for $o \in A \setminus L$ yields the top-ranked alternative of each agent as her representative. Any other lottery y' must have $y'(o) < 1/\ell$ for some alternative $o \in L$, implying that the agent who has alternative o as her top choice prefers lottery y to y' . Thus, lottery y dominates every other lottery (i.e., it is the unique efficient lottery for the quantile parameter vector $\mathbf{h}^2$), including the set of (deterministic) lotteries that are efficient for the quantile parameter vector $\mathbf{h}^1$. We conclude that, if $|L| \geq 2$, no lottery is efficient for both quantile parameter vectors $\mathbf{h}^1$ and $\mathbf{h}^2$, completing the “only if” part of the proof. $\square$

A (randomized) voting rule² R takes as input a profile and returns a lottery over the alternatives in A . Given a profile $\succ$ and an alternative $a \in A$, we denote by $R_a(\succ)$ the probability assigned to alternative a when the voting rule R is applied on the profile $\succ$. A randomized voting rule is efficient if it produces an efficient lottery. Next, we define strategyproofness for voters with quantile utilities.

Definition 5 (strategyproofness). *Given a set N of n agents with quantile parameter vector $\mathbf{h}$ and a set A of m alternatives, the voting rule R is strategyproof for the domain $(N, A, \cdot, \mathbf{h})$ if for every preference profile $\succ$, agent i , and preference $\succ'_i$, it holds $\text{rep}(R(\succ), \succ_i, h_i) \succeq_i \text{rep}(R(\succ_{-i}, \succ'_i), \succ_i, h_i)$.*

As usual, the notation $(\succ_{-i}, \succ'_i)$ is used to denote the profile that is obtained from $\succ$ when agent i changes her preference to $\succ'_i$.

A natural task is to characterize strategyproof voting rules for agents with quantile utilities. Towards this, we first define monotonic voting rules for voting instances with two alternatives.

²Randomized voting rules are also called *social decision schemes* or *probabilistic social choice functions* in the literature; e.g., see [14].

Definition 6 (monotonicity). Consider the voting rule R applied on profiles consisting of the preferences of a set N of agents and a set $A = \{a, b\}$ of two alternatives. The rule R is monotonic if $R_a(\succ) \geq R_a(\succ_{-i}, \succ'_i)$ for every preference profile $\succ$, every agent $i \in N$ with $a \succ_i b$, and preference $\succ'_i$ with $b \succ'_i a$.

We show the relation between monotonic voting rules and strategyproofness.

Theorem 7. A monotonic voting rule for two alternatives is strategyproof for any quantile parameter vector. For any non-monotonic rule R , there exists a quantile parameter vector $\mathbf{h}$ such that R is not strategyproof.

Proof. Consider the monotonic voting rule R applied on a preference profile $\succ$ with two alternatives a and b . For each agent $i \in N$ with $a \succ_i b$, it suffices to consider the case where $\text{rep}(R(\succ), \succ_i, h_i) = b$ and show that $\text{rep}(R(\succ_{-i}, \succ'_i), \succ_i, h_i) = b$ as well, where $b \succ'_i a$. Assume otherwise that $\text{rep}(R(\succ_{-i}, \succ'_i), \succ_i, h_i) = a$. By the definition of $\text{rep}(R(\succ), \succ_i, h_i)$, we have that $R_b(\succ) > h_i$, i.e., $R_a(\succ) < 1 - h_i$. By the definition of $\text{rep}(R(\succ_{-i}, \succ'_i), \succ_i, h_i)$, we have that $R_b(\succ_{-i}, \succ'_i) \leq h_i$, i.e., $R_a(\succ_{-i}, \succ'_i) \geq 1 - h_i$. Thus, $R_a(\succ) < R_a(\succ_{-i}, \succ'_i)$, contradicting the monotonicity of R .

Consider a non-monotonic voting rule R , a preference profile $\succ$, an agent i with $a \succ_i b$, and another preference $\succ'_i$ with $b \succ'_i a$ such that $t_1 = R_a(\succ) < R_a(\succ_{-i}, \succ'_i) = t_2$. By setting $h_i = 1 - \frac{t_1+t_2}{2}$ (notice that this definition yields a valid quantile parameter with $0 < h_i < 1$), we have $\text{rep}(R(\succ), \succ_i, h_i) = b$ and $\text{rep}(R(\succ_{-i}, \succ'_i), \succ_i, h_i) = a$, which violates the strategyproofness of R . Indeed, $R_b(\succ) = 1 - R_a(\succ) = 1 - t_1 > 1 - \frac{t_1+t_2}{2} = h_i$ (hence, $\text{rep}(R(\succ), \succ_i, h_i) = b$), and $R_b(\succ_{-i}, \succ'_i) = 1 - R_a(\succ_{-i}, \succ'_i) = 1 - t_2 \leq 1 - \frac{t_1+t_2}{2} = h_i$ and $R_a(\succ_{-i}, \succ'_i) = t_2 > t_1 \geq 0$ (hence, $\text{rep}(R(\succ_{-i}, \succ'_i), \succ_i, h_i) = a$). $\square$

When efficiency and strategyproofness are defined using stochastic dominance, the celebrated result of Gibbard [24] states that random dictatorships are the only strategyproof and efficient³ voting rules. In the remainder of this section, we explore whether analogous characterizations for strategyproof and efficient rules exist in our model (under Definitions 3 and 5).

For two alternatives, we define the randomized voting rule Q -plurality (or Q for short) as follows. Consider the voting instance $(N, \{a, b\}, \succ, \mathbf{h})$ with two alternatives $\{a, b\}$ and, without loss of generality, assume that the plurality winner⁴ in $\succ$ is alternative a . Define $Q_a(\succ) = 1 - \min_{i:a \succ_i b} h_i$ (and $Q_b(\succ) = \min_{i:a \succ_i b} h_i$).

Theorem 8. Q -plurality is an efficient and strategyproof rule for voting instances with two alternatives.

Proof. Consider a voting instance $(N, \{a, b\}, \succ, \mathbf{h})$ in which (without loss of generality) alternative a is the plurality winner. Denote by $W = \{i \in N : a \succ_i b\}$ the set of agents who prefer alternative a to alternative b . Notice that $\text{rep}(Q(\succ), \succ_j, h_j) = a$ for every agent $j \in W$; indeed, $Q_b(\succ) = \min_{i \in W} h_i \leq h_j$ and $Q_a(\succ) = 1 - \min_{i \in W} h_i > 0$. Let $L = \{i \in N : b \succ_i a, \text{rep}(Q(\succ), \succ_i, h_i) = a\}$ be the set of agents who have their least preferred alternative as representative. We will prove the efficiency of Q -plurality by showing that any lottery that makes alternative b the representative for some agent in L (and is, thus, more preferable than Q for this agent) makes the least preferred alternative the representative for some agent in W (and is, thus, less preferable than Q for this agent).

Consider the lottery x that satisfies $\text{rep}(x, \succ_j, h_j) = b$ for some agent $j \in L$ and let i^* be an agent in W with minimum quantile parameter, i.e., $h_{i^*} = \min_{i \in W} h_i$. Since $\text{rep}(x, \succ_j, h_j) = b$, $\text{rep}(Q(\succ), \succ_j, h_j) = a$, and by the definition of lottery Q and the quantile parameter h_{i^*} , we have

$$x(a) \leq h_j < Q_a(\succ) = 1 - \min_{i \in W} h_i = 1 - h_{i^*}.$$

³Efficiency, here, is equivalent to *ex post* efficiency. A voting rule is a random dictatorship if it returns the top alternative of a randomly selected agent, according to some probability distribution.

⁴We assume that ties are always broken in favor of the same alternatives.

Equivalently, $x(b) = 1 - x(a) > h_{i^*}$, which implies that $\text{rep}(x, \succ_{i^*}, h_{i^*}) = b$, as desired.

To prove strategyproofness, it suffices to consider an agent $j \in L$ and show that alternative a is still her representative under Q -plurality after she changes her preference to $a \succ'_j b$. Clearly, alternative a is still the plurality winner in profile $(\succ_{-j}, \succ'_j)$ and, hence,

$$Q_a(\succ_{-j}, \succ'_j) = 1 - \min_{i \in W \cup \{j\}} h_i \geq 1 - \min_{i \in W} h_i = Q_a(\succ) > h_j,$$

which implies that $\text{rep}(Q(\succ_{-j}, \succ'_j), \succ_j, h_j) = a$, as desired. The last inequality follows since $j \in L$ and, hence, $\text{rep}(Q(\succ), \succ_j, h_j) = a$. $\square$

For more than two alternatives, the picture changes. Our next statement demonstrates that, rather unsurprisingly, deterministic dictatorships are the only efficient and strategyproof voting rules when the agents have small quantile parameters. However, positive results are possible for quantile parameters higher than $1/2$.

Theorem 9. *Consider a set N of n agents with the quantile parameter vector $\mathbf{h}$ and a set of alternatives $A = \{a, b, c\}$.*

- (a) *If $\mathbf{h} \in [0, 1/3]^n$, any efficient and strategyproof voting rule is a deterministic dictatorship.*
- (b) *If $\mathbf{h} \in [1/2, 2/3]^n$, the voting rule that returns each of the two alternatives with the highest plurality score with probability $1/2$ is efficient and strategyproof.*
- (c) *If $\mathbf{h} \in [2/3, 1]^n$, the voting rule that returns each alternative with probability $1/3$ is efficient and strategyproof.*

Proof. **(a)** Consider any lottery x over the alternatives in A assume that $x(a) = \max_{o \in A} x(o)$. Then, $x(a) \geq 1/3$ and $\sum_{o: a \succeq_i o} x(o) \geq 1/3 > h_i$ for every $i \in N$. This implies that $a \succeq_i \text{rep}(x, \succ_i, h_i)$, meaning that the lottery x is either dominated by the deterministic lottery $y(a) = 1$ (and $y(b) = y(c) = 0$) or the two lotteries are equivalent in the sense that, for any agent, they yield the same alternative as representative. Hence, any efficient voting rule can be considered as deterministic in this case, and the statement follows from the Gibbard-Satterthwaite theorem [23, 40].

(b) Consider a profile $\succ$ and let a and b be the two alternatives with the highest plurality score. Then, the lottery x with $x(a) = x(b) = 1/2$ and $x(c) = 0$ yields as representative in each agent her highest-ranked alternative among a and b .

We first show that x is efficient. Notice that lottery x can be dominated only by some lottery y that makes alternative c the representative of an agent j who has alternative c ranked first. By the definition of the representative $\text{rep}(y, \succ_j, h_j)$ and the assumption on h_j , it must be $y(a) + y(b) \leq h_j < 2/3$. Thus, lottery y must assign probability smaller than $1/3$ to at least one of the alternatives a and b . Assume, without loss of generality, that $y(a) < 1/3$ and consider an agent i with alternative a as her top choice (such an agent exists since both alternatives a and b have at least as high plurality score as alternative c). Then, $y(b) + y(c) = 1 - y(a) > 2/3 > h_i$, meaning that $a \succ_i \text{rep}(y, \succ_i, h_i)$. Hence, lottery y does not dominate lottery x and lottery x is efficient.

To prove strategyproofness, it suffices to consider an agent $j \in N$ with c as top-ranked alternative and, thus, her second-ranked alternative as representative under lottery x . Now, observe that any deviating preference $\succ'_j$ by the agent $j \in N$ cannot increase the plurality score of c or decrease the plurality score of a and b in $(\succ_{-j}, \succ'_j)$ compared to $\succ$. Thus, a and b are still the two alternatives with the highest plurality score, and the lottery returned by the rule is the same.

(c) In this case, the uniform lottery with $x(a) = x(b) = x(c) = 1/3$ makes the top-ranked alternative of each agent her representative. Hence, this rule is trivially efficient and strategyproof. $\square$

With our Theorem 9, we have made some initial progress towards answering the following challenging open problem.

Open Problem 1. *Characterize the efficient and strategyproof randomized voting rules for voting settings with at least three alternatives, any number of agents, and all quantile parameter vectors.*

4 One-sided matching

In this section, we consider settings with a set N of n agents and a set M of n items. Each agent $i \in N$ has a quantile parameter $h_i \in [0, 1)$ and a preference order $\succ_i$ over the items in M . A matching mechanism takes as input a preference profile $\succ$ and the vector $\mathbf{h}$ of quantile parameters, and returns a lottery over perfect matchings between the sets N and M . Thus, a lottery $\mathbf{x}$ assigns probability x_{ig} for agent $i \in N$ and item $g \in M$ such that $\sum_{i \in N} x_{ig} = 1$ for each $g \in M$ and $\sum_{g \in M} x_{ig} = 1$ for each $i \in N$. We usually view the lottery $\mathbf{x}$ over matchings as n lotteries over the items in M , with one lottery x_i for each agent $i \in N$.

We begin our study of lotteries over matchings by focusing on the maximization of (a variant of) social welfare. The proof of the next statement appears in Appendix B.

Theorem 10. *Given a one-sided matching instance $(N, M, \succ, \mathbf{h})$, a lottery that maximizes the number of agents having their top-ranked item as representative can be computed in polynomial time.*

One can consider other social welfare objectives for lotteries over matchings. We believe that the following two problems are worth studying.

Open Problem 2. *What is the complexity of the following two problems?*

- *Given a one-sided matching instance $(N, M, \succ, \mathbf{h})$ and an integer rank requirement r_i for every agent $i \in N$, compute a lottery $\mathbf{x}$ that maximizes $|\{i \in N : \text{rank}_i(\text{rep}(x_i, \succ_i, h_i)) \in [r_i]\}|$, i.e., the number of agents with a representative meeting their rank requirement. The problem solved in Theorem 10 is the special case with $r_i = 1$ for every agent $i \in N$.*
- *Given a one-sided matching instance $(N, M, \succ, \mathbf{h})$, compute a lottery $\mathbf{x}$ that minimizes $\sum_{i \in N} \text{rank}_i(\text{rep}(x_i, \succ_i, h_i))$, i.e., the sum of ranks of the representatives over all agents.*

We now adapt the definitions of efficiency and strategyproofness for lotteries over matchings and matching mechanisms.

Definition 11 (efficiency for lotteries over matchings). *Given a one-sided matching instance $(N, M, \succ, \mathbf{h})$, a lottery $\mathbf{x}$ over matchings is efficient if there is no other lottery $\mathbf{y}$ such that $\text{rep}(y_i, \succ_i, h_i) \succeq_i \text{rep}(x_i, \succ_i, h_i)$ for every agent $i \in N$ and there is an agent $i^* \in N$ such that $\text{rep}(y_{i^*}, \succ_{i^*}, h_{i^*}) \succ_{i^*} \text{rep}(x_{i^*}, \succ_{i^*}, h_{i^*})$.*

Definition 12 (strategyproofness of matching mechanisms). *Given a set N of n agents with quantile parameter vector $\mathbf{h}$ and a set M of n items, the matching mechanism R is strategyproof for the domain $(N, M, \cdot, \mathbf{h})$ if for every preference profile $\succ$, agent $i \in N$, and preference $\succ'_i$, it holds $\text{rep}(R(\succ), \succ_i, h_i) \succeq_i \text{rep}(R(\succ_{-i}, \succ'_i), \succ_i, h_i)$.*

We will prove that efficiency and strategyproofness are compatible for agents with quantile utilities. The main component of our proof is a simple set of linear inequalities that check whether lotteries over matchings that satisfy given requirements for the ranks of the representatives of the agents exist. For a given one-sided matching instance $(N, M, \succ, \mathbf{h})$ and a vector of rank requirements $\mathbf{r} = (r_1, r_2, \dots, r_n)$

with integer $r_i \in [n]$, the following linear program is feasible if and only if there is a lottery over matchings that satisfy the rank requirement r_i for the representative of each agent $i \in N$.

$$\begin{aligned} \sum_{g: \text{rank}_i(g) \leq r_i} x_{ig} &\geq 1 - h_i, \quad \forall i \in N \\ \sum_{g \in M} x_{ig} &= 1, \quad \forall i \in N \\ \sum_{i \in N} x_{ig} &= 1, \quad \forall g \in M \end{aligned}$$

Dropping N , M , and $\mathbf{h}$ from notation, we will refer to this linear program as $\text{LP}(\succ; \mathbf{r})$.

The compatibility of efficiency and strategyproofness can be obtained via the following variant of the *serial dictatorship* mechanism.

Serial dictatorship (SD) mechanism: On input a one-sided matching instance $(N, M, \succ, \mathbf{h})$ with n agents/items, mechanism SD starts with a vector $\mathbf{r} = (n, \dots, n)$ of minimum rank requirements for all agents and considers the agents one by one in increasing order of their ids. When the agent $i \in N$ is considered, the mechanism computes the minimum rank t for agent i so that the linear program $\text{LP}(\succ; \mathbf{r}_{-i}, t)$ is feasible and updates $\mathbf{r}$ by setting $r_i = t$. After all agents have been considered, an arbitrary lottery over matchings that satisfies $\text{LP}(\succ; \mathbf{r})$ is returned as the output of the mechanism.

Theorem 13. *Mechanism SD is efficient and strategyproof.*

Proof. The crucial property of mechanism SD is that the final linear program $\text{LP}(\succ; \mathbf{r})$ has as solutions all those lotteries which make representative for agent 1 her top-ranked item and representative for agent $i \geq 2$ her highest-ranked item under rank-constraints for the representative items for agents $1, \dots, i - 1$.

To prove efficiency, let $\mathbf{x}$ be the lottery returned by the mechanism as the solution of the linear program $\text{LP}(\succ; \mathbf{r})$. For the sake of contradiction, assume that there is another lottery $\mathbf{y}$ such that $\text{rank}_i(y_i, \succ_i, h_i) \leq \text{rank}_i(x_i, \succ_i, h_i)$ for every agent $i \in N$ and $\text{rank}_{i^*}(y_{i^*}, \succ_{i^*}, h_{i^*}) < \text{rank}_{i^*}(x_{i^*}, \succ_{i^*}, h_{i^*})$ for some agent $i^* \in N$. Let $i' \in N$ be the agent of minimum id with $r'_{i'} = \text{rank}_{i'}(y_{i'}, \succ_{i'}, h_{i'}) < r_i$, i.e., $\text{rank}_i(y_i, \succ_i, h_i) = r_i$ for $i = 1, \dots, i' - 1$. This means that the linear program $\text{LP}(\succ; (r_1, \dots, r_{i'-1}, r'_{i'}, n, \dots, n))$ has lottery $\mathbf{y}$ as solution, contradicting the definition of $r_{i'}$ at round i' of the mechanism.

Strategyproofness follows since when it considers agent $i \in N$, the mechanism restricts the candidate output lotteries to those that make the best possible item a representative for agent $i \in N$, given decisions about the representatives of agents $1, \dots, i - 1$ in previous rounds. Hence, unilateral misreporting by agent i cannot result to a final lottery that improves her representative item further. $\square$

In the following, we define two fairness properties for lotteries over matchings and study their interplay with efficiency (and strategyproofness). We begin with the definition of *proportionality*.

Definition 14 (proportionality). *A lottery $\mathbf{x}$ for the one-sided matching instance $(N, M, \succ, \mathbf{h})$ is proportional if $\text{rank}_i(\text{rep}(x_j, \succ_i, h_i)) \leq \lceil n(1 - h_i) \rceil$ for every agent $i \in N$.*

Note that our definition implies that the uniform lottery is proportional. The first question that arises is whether proportionality is compatible with efficiency. We prove that this is indeed the case, using a straightforward restriction of the serial dictatorship mechanism. In its definition below, we use p_i as a shorthand of the rank requirement for agent $i \in N$ from Definition 14, i.e., $p_i = \lceil n(1 - h_i) \rceil$.

Proportionality-constraint serial dictatorship (PSD): On input a one-sided matching instance $(N, M, \succ, \mathbf{h})$ with n agents/items, mechanism PSD starts with a vector $\mathbf{r} = (p_1, \dots, p_n)$ of the rank requirements for proportionality for all agents and considers the agents one by one in increasing order of their ids. When the agent $i \in N$ is considered, the mechanism computes the minimum rank t for agent i so that the linear program $\text{LP}(\succ; \mathbf{r}_{-i}, t)$ is feasible and updates $\mathbf{r}$ by setting $r_i = t$. After all agents have been considered, an arbitrary lottery over matchings that satisfies $\text{LP}(\succ; \mathbf{r})$ is returned as the output of the mechanism. The proof of the next statement appears in Appendix B.

Theorem 15. *Mechanism PSD returns efficient and proportional lotteries.*

Due to its similarity with mechanism SD, it is tempting to assume that the PSD mechanism is also strategyproof. Surprisingly, this is not the case; the restriction of serial dictatorships to proportional lotteries violates strategyproofness.

Theorem 16. *Mechanism PSD is not strategyproof.*

Proof. Consider a one-sided matching instance with three agents 1, 2, and 3 and three items a, b , and c . Agent 1 has $h_1 = 0$ and preference $a \succ_1 b \succ_1 c$. Agents 2 and 3 have $h_2 = h_3 = 1/3$ and identical preference $\succ_2 = \succ_3$ with $b \succ_2 c \succ_2 a$.

Proportionality restricts the outcome of mechanism PSD to lotteries in which some of the items b and c are representative of agents 2 and 3, while any item can be the representative of agent 1. Under this restriction, the lottery selected by the mechanism yields the representatives a, b , and c for the three agents, respectively. E.g., such a lottery has $x_{1a} = 1, x_{2b} = x_{3c} = 2/3, x_{2c} = x_{3b} = 1/3$, and $x_{1b} = x_{1c} = x_{2a} = x_{3a} = 0$. Notice that making item b representative in both agents 2 and 3 is infeasible since both x_{2b} and x_{3b} should be at least $2/3$ violating the matching condition.

Now, assume that agent 3 misreports the preference $a \succ'_3 b \succ'_3 c$. Proportionality restricts the outcome of mechanism PSD to lotteries in which the representative of agent 2 is some of the items b and c and the representative of agent 3 is some of the items a and b . Among them, the set of lotteries that make item a representative of agent 1 is non-empty, and the matching and proportionality conditions yield that items c and b are the representatives of agents 2 and 3, respectively. Indeed, such a lottery $\mathbf{x}'$ should satisfy $x'_{1a} = 1$ and, by the matching condition, $x'_{3a} = 0$. Hence, item a cannot be the representative of agent 3. To make item b her representative as proportionality requires, it must be $x'_{3b} \geq 2/3$. But then, it must also be $x'_{2b} \leq 1/3$, which implies that item b cannot be the representative of agent 2. The lottery $\mathbf{x}'$ with $x'_{1a} = 1, x'_{2b} = x'_{3c} = 1/3, x'_{2c} = x'_{3b} = 2/3$, and $x'_{1b} = x'_{1c} = x'_{2a} = x'_{3a} = 0$ makes items a, c , and b representatives in agents 1, 2, and 3, respectively.

Thus, when agent 3 misreports $\succ'_3$ instead of her true preference $\succ_2$, mechanism PSD returns a lottery with a strictly better representative for her, violating strategyproofness. $\square$

Now, Theorems 15 and 16 beg the following question.

Open Problem 3. *Are there efficient, strategyproof, and proportional mechanisms for one-sided matching instances with quantile agent utilities?*

We now define *envy-freeness*, our second fairness property.

Definition 17 (envy-freeness). *A lottery x for the one-sided matching instance $(N, M, \succ, \mathbf{h})$ is envy-free if $\text{rep}(x_i, \succ_i, h_i) \succeq_i \text{rep}(x_j, \succ_i, h_i)$ for every pair of agents $i, j \in N$.*

The definition of envy-freeness is a very natural one. One may wonder why the particular definition of proportionality is well justified. The next statement indicates that, similar to the relations of these concepts in the literature on fair division, our definitions are such that envy-freeness implies proportionality.

Theorem 18. *Envy-freeness implies proportionality.*

Proof. Consider a matching instance and let $\mathbf{x}$ be an envy-free lottery over matchings. Abbreviate the rank requirement for agent $i \in N$ for proportionality by $p_i = \lceil n(1 - h_i) \rceil$ and assume that $\mathbf{x}$ is not proportional. Then, there exists an agent $i^* \in N$ such that

$$\text{rank}_{i^*}(\text{rep}(x_{i^*}, \succ_{i^*}, h_{i^*})) \geq p_{i^*} + 1.$$

Since $\mathbf{x}$ is envy-free, we have

$$\text{rank}_{i^*}(\text{rep}(x_i, \succ_{i^*}, h_{i^*})) \geq \text{rank}_{i^*}(\text{rep}(x_{i^*}, \succ_{i^*}, h_{i^*})) \geq p_{i^*} + 1, \quad (6)$$

for every agent $i \in N$. Denote by g^* the item ranked p_{i^*} -th by agent i^* . By the definition of $\text{rep}(x_i, \succ_{i^*}, h_{i^*})$ and inequality (6), we have

$$\sum_{g \in M: g^* \succ_{i^*} g} x_{ig} > h_{i^*}$$

for every agent $i \in N$. Using this inequality, we have

$$\sum_{g \in M: g^* \succ_{i^*} g} \sum_{i \in N} x_{ig} = \sum_{i \in N} \sum_{g \in M: g^* \succ_{i^*} g} x_{ig} > n \cdot h_{i^*}.$$

Now, notice that there are $n - p_{i^*} = n - \lceil n(1 - h_{i^*}) \rceil \leq n \cdot h_{i^*}$ items ranked below the item g^* by agent i^* and, hence, at most $n \cdot h_{i^*}$ terms in the outer sum of the LHS of the last equation. Thus, there exists one such term corresponding to an item $g' \in M$ such that $\sum_{i \in N} x_{ig'} > 1$, violating the validity of lottery $\mathbf{x}$. $\square$

Unfortunately, envy-freeness and efficiency may not be possible simultaneously, as the next simple example illustrates.

Example. Consider an instance with two agents 1 and 2 with identical preference $a \succ b$ for two items a and b and identical quantile parameter $h < 1/2$. First, observe that the same item cannot become the representative of both agents. If item a is the representative for both agents, this means that $x_{1a} \geq 1 - h > 1/2$ and $x_{2a} \geq 1 - h > 1/2$, violating the matching constraint $x_{1a} + x_{2a} \leq 1$. The lottery that makes item b representative for both agents (e.g., the uniform lottery) is dominated by the lotteries that make item a representative for agent 1 and item b representative for agent 2 (e.g., $x_{1a} = x_{2b} = 1$ and $x_{1b} = x_{2a} = 0$) and vice-versa. These two classes of efficient lotteries are clearly not envy-free. $\square$

The following question is, hence, interesting to study.

Open Problem 4. *What is the complexity of deciding whether an efficient and envy-free lottery exists for a given one-sided matching instance with quantile agent utilities?*

5 Conclusion

We have initiated a novel direction in randomized social choice by proposing the new concept of quantile agent utility for randomized outcomes. We presented a series of results that demonstrate the power of the new definition in various social choice settings, namely, voting, one-sided matchings, and stable matchings, as well as its implications for concepts such as efficiency, strategyproofness, and fairness. We have identified several open problems, suggesting interesting future directions. There are other scopes for future works as well. For example, we can ask for mechanisms that are strategyproof even when an agent can misreport their preference and quantile parameter. Such questions deserve further investigation.

Acknowledgements

We thank Felix Brandt and Herve Moulin for helpful discussions and pointers to the literature. This research was initiated at the Workshop on Multi-Agent Systems held at Bellairs Research Institute, Barbados, in March 2023. The work of Caragiannis has been partially supported by Independent Research Fund Denmark (DFF) under grant 2032-00185B.

References

- [1] Atila Abdulkadiroğlu and Tayfun Sönmez. Ordinal efficiency and dominated sets of assignments. *Journal of Economic Theory*, 112(1):157–172, 2003.
- [2] Atila Abdulkadiroğlu and Tayfun Sönmez. School choice: A mechanism design approach. *American Economic Review*, 93(3):729–747, 2003.
- [3] Hernán G. Abeledo and Uriel G. Rothblum. Stable matchings and linear inequalities. *Discrete Applied Mathematics*, 54(1):1–27, 1994.
- [4] Ron Aharoni and Tamás Fleiner. On a lemma of Scarf. *Journal of Combinatorial Theory, Series B*, 87:72–80, 2003.
- [5] Ahmet Alkan and David Gale. Stable schedule matching under revealed preference. *Journal of Economic Theory*, 112(2):289–306, 2003.
- [6] Kenneth J. Arrow. *Social Choice and Individual Values*. John Wiley & Sons, 2nd edition, 1963 edition, 1951.
- [7] Haris Aziz. A probabilistic approach to voting, allocation, matching, and coalition formation. In Jean-François Laslier, Hervé Moulin, M. Remzi Sanver, and William S. Zwicker, editors, *The Future of Economic Design*, pages 45–50. Springer, 2019.
- [8] Haris Aziz and Bettina Klaus. Random matching under priorities: Stability and no envy concepts. *Social Choice and Welfare*, 53(2):213–259, 2019.
- [9] Haris Aziz, Florian Brandl, Felix Brandt, and Markus Brill. On the tradeoff between efficiency and strategyproofness. *Games and Economic Behavior*, 110:1–18, 2018.
- [10] Haris Aziz, Rupert Freeman, Nisarg Shah, and Rohit Vaish. Best of both worlds: Ex ante and ex post fairness in resource allocation. *Operations Research*, 72(4):1674–1688, 2024.
- [11] Moshe Babaioff, Tomer Ezra, and Uriel Feige. On best-of-both-worlds fair-share allocations. In *Proceedings of the 18th International Conference on Web and Internet Economics (WINE)*, pages 237–255, 2022.
- [12] Péter Biró, Katarina Cechlárová, and Tamás Fleiner. The dynamics of stable matchings and half-matchings for the stable marriage and roommates problems. *International Journal of Game Theory*, 36(3):333–352, 2008.
- [13] Anna Bogomolnaia and Hervé Moulin. A new solution to the random assignment problem. *Journal of Economic Theory*, 100(2):295–328, 2001.
- [14] Felix Brandt. Rolling the dice: Recent results in probabilistic social choice. In Ulle Endriss, editor, *Trends in Computational Social Choice*, pages 3–26. AI Access, 2017.
- [15] Ioannis Caragiannis, Aris Filos-Ratsikas, Panagiotis Kanellopoulos, and Rohit Vaish. Stable fractional matchings. *Artificial Intelligence*, 295:103416, 2020.
- [16] Ioannis Caragiannis, Panagiotis Kanellopoulos, and Maria Kyropoulou. On interim envy-free allocation lotteries. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*, pages 264–284, 2021.
- [17] Jiehua Chen, Sanjukta Roy, and Manuel Sorge. Fractional matchings under preferences: Stability and optimality. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 89–95, 2021.
- [18] Battal Doğan and Kemal Yıldız. Efficiency and stability of probabilistic assignments in marriage problems. *Games and Economic Behavior*, 95:47–58, 2016.

- [19] John Duggan and Thomas Schwartz. Strategic manipulability without resoluteness or shared beliefs: Gibbard-satterthwaite generalized. *Social Choice and Welfare*, 17(1):85–93, 2000.
- [20] Peter C. Fishburn. Lotteries and social choices. *Journal of Economic Theory*, 5(2):189–207, 1972.
- [21] Duncan Karl Foley. *Resource allocation and the public sector*. Yale University, 1966.
- [22] David Gale and Lloyd S. Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 120(5):386–391, 1962.
- [23] Allan Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41(4):587–602, 1973.
- [24] Allan Gibbard. Manipulation of schemes that mix voting with chance. *Econometrica*, 45(3):665–681, 1977.
- [25] Dan Gusfield and Robert W. Irving. *The Stable Marriage Problem—Structure and Algorithms*. Foundations of Computing Series. MIT Press, 1989.
- [26] Roger A. Horn and Charles R. Johnson. *Topics in Matrix Analysis*. Cambridge University Press, 1991.
- [27] Aanund Hylland and Richard Zeckhauser. The efficient allocation of individuals to positions. *Journal of Political Economy*, 87(2):293–314, 1979.
- [28] Michael D. Intriligator. A probabilistic model of social choice. *The Review of Economic Studies*, 40(4):553–560, 1973.
- [29] Takashi Ishizuka and Naoyuki Kamiyama. On the complexity of stable fractional hypergraph matching. In *Proceedings of the 29th International Symposium on Algorithms and Computation (ISAAC)*, volume 123 of *LIPIcs*, pages 11:1–11:12, 2018.
- [30] Onur Kesten and M. Utku Ünver. A theory of school-choice lotteries. *Theoretical Economics*, 10(2): 543–595, 2015.
- [31] Shiva Kintali, Laura J. Poplawski, Rajmohan Rajaraman, Ravi Sundaram, and Shang-Hua Teng. Reducibility among fractional stability problems. *SIAM Journal on Computing*, 42(6):2063–2113, 2013.
- [32] Harold W. Kuhn. The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2(1-2):83–97, 1955.
- [33] Vikram Manjunath. Stability and the core of probabilistic marriage problems. *Available at SSRN*, 2013.
- [34] Vikram Manjunath. Markets for fractional partnerships. *Available at SSRN*, 2014.
- [35] David G. McVitie and Leslie B. Wilson. The stable marriage problem. *Communications of the ACM*, 14(7):486–490, 1971.
- [36] Ariel D. Procaccia. Can approximation circumvent Gibbard-Satterthwaite? In *Proceedings of the 26th Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 597–606, 2010.
- [37] Alvin E. Roth. The economics of matching: Stability and incentives. *Mathematics of Operations Research*, 7(4):617–628, 1982.
- [38] Alvin E. Roth, Uriel G. Rothblum, and John H. Vande Vate. Stable matchings, optimal assignments, and linear programming. *Mathematics of Operations Research*, 18(4):803–828, 1993.
- [39] Shin Sato. On strategy-proof social choice correspondences. *Social Choice and Welfare*, 31(2): 331–343, 2008.
- [40] Mark Allen Satterthwaite. Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2):187–217, 1975.
- [41] Hugo Steinhaus. The problem of fair division. *Econometrica*, 16:101–104, 1948.
- [42] Chung-Piaw Teo and Jay Sethuraman. The geometry of fractional stable matchings and its applications. *Mathematics of Operations Research*, 23(4):874–891, 1998.
- [43] John H. Vande Vate. Linear programming brings marital bliss. *Operations Research Letters*, 8(3): 147–153, 1989.
- [44] Hal R. Varian. Equity, envy, and efficiency. *Journal of Economic Theory*, 9(1):63–91, 1974.
- [45] Richard Zeckhauser. Majority rule with lotteries on alternatives. *The Quarterly Journal of Economics*,

83(4):696–703, 1969.

Ioannis Caragiannis
Aarhus University
Aarhus, Denmark
Email: iannis@cs.au.dk

Sanjukta Roy
Indian Statistical Institute
Kolkata, India
Email: sanjukta@isical.ac.in

A More related work

In the two-sided matching literature, a version of the random-matching problem, where both sides of the market have strict preferences, was studied by Roth et al. [38], who showed that the set of fractional stable matchings has a lattice structure. One can interpret a fractional deterministic matching as a random matching using Birkhoff-von Neumann's theorem [26, Theorem 3.2.6]. Thus, stable fractional matchings has been studied as a relaxation of the integrality constraint in the linear equations encoding stability [43, 3, 42] which were forerunner of numerous other works [12, 31, 30, 29, 8, 15, 18, 33, 34]. In particular, stable fractional matchings have been considered by Aharoni and Fleiner [4], who proposed the definition of ordinal stability. We remark that, in our setting with agents having quantile utilities, stability generalizes both ordinal and integral stability. Aharoni and Fleiner [4] studied ordinal stability in the hypergraphic setting and found that Scarf's lemma from game theory guarantees the existence of ordinally stable matchings. Later, Aziz and Klaus [8] studied multiple notions of stability and the relations between them. Alkan and Gale [5] studied a type of stability for fractional matching when both sides are equipped with complex preference structures. Caragiannis et al. [15] defined cardinal stable lotteries when both sides have cardinal preferences as opposed to rankings or ordinal preferences. Chen et al. [17] studied both cardinal and ordinal stability. Lotteries have been extensively studied for another application of matching under preferences, namely, the school choice problem [2, 30]. In particular, Kesten and Ünver [30] studied ordinal stability, which they refer to as ex-ante stability.

B Omitted proofs from Section 4

Proof of Theorem 10. For each item $g \in M$, denote by S_g the set of agents having item g as their top choice. An agent from S_g can have item g as representative only if the lottery assigns her to item g with probability at least $1 - h_i$. Denote by $\bar{S}_g$ a subset of S_g of maximum cardinality that satisfies $\sum_{i \in \bar{S}_g} (1 - h_i) \leq 1$; this condition is necessary so that there is a lottery that makes item g representative for all agents in $\bar{S}_g$. Then, the maximum number of agents having their top item as representative is at most $\sum_{g \in M} |\bar{S}_g|$.

We now present a lottery under which a maximum number of $\sum_{g \in M} |\bar{S}_g|$ agents have their top item as representative. Notice that for every two different items $g_1, g_2 \in M$, the sets $\bar{S}_{g_1}$ and $\bar{S}_{g_2}$ are disjoint. Thus, by setting $x_{ig} = 1 - h_i$ for every $g \in M$ and $i \in \bar{S}_g$, we get that all agents in $\cup_{g \in M} \bar{S}_g$ have their top item as representative. By the definition of set $\bar{S}_g$, we have the condition $\sum_{i \in \bar{S}_g} x_{ig} \leq 1$ for every item g . Hence, we can trivially complete $\mathbf{x}$ and get a valid lottery by setting the values x_{ig} for every other agent-item pair (i, g) with $i \notin \bar{S}_g$ so that $\sum_{i \in N} x_{ig} = 1$ for every item $g \in M$ and $\sum_{g \in M} x_{ig} = 1$ for every agent $i \in N$.

Notice that, for a given item $g \in M$, the set $\bar{S}_g$ can be easily computed by starting with the empty set, considering the agents in set S_g in monotone non-increasing order of their quantile parameter, and including an agent in $\bar{S}_g$ as long as the sum of the quantities $1 - h_i$ in $\bar{S}_g$ does not exceed 1. $\square$

Proof of Theorem 15. The crucial property of mechanism PSD is that the final linear program $\text{LP}(\succ; \mathbf{r})$ has as solutions all those lotteries which make representative for agent 1 her top-ranked item and representative for agent $i \geq 2$ her highest-ranked item under rank-constraints for the representative items for agents $1, \dots, i - 1$ and proportionality rank-constraints for agents $i + 1, \dots, n$.

The proportionality of the final lottery is guaranteed by the definition of the mechanism. Initially, before the execution of the first round, the linear program $\text{LP}(\succ, p_1, \dots, p_n)$ is feasible (e.g., it is satisfied by the uniform lottery). When considering agent $i \in N$, the linear program $\text{LP}(\succ; (r_1, r_2, \dots, r_{i-1}, p_i, \dots, p_n))$ is guaranteed to be feasible after the execution of phase $i - 1$. Hence, $r_i \geq p_i$ and $\text{LP}(\succ; \mathbf{r})$ has proportional lotteries as feasible solutions.

Efficiency follows by slightly adapting the argument we used in the proof of Theorem 13; we include it here for completeness. To prove efficiency, let $\mathbf{x}$ be the lottery returned by the mechanism as the solution of the linear program $\text{LP}(\succ; \mathbf{r})$. For the sake of contradiction, assume that there is another lottery $\mathbf{y}$ such that $\text{rank}_i(y_i, \succ_i, h_i) \leq \text{rank}_i(x_i, \succ_i, h_i)$ for every agent $i \in N$ and $\text{rank}_{i^*}(y_{i^*}, \succ_{i^*}, h_{i^*}) < \text{rank}_{i^*}(x_{i^*}, \succ_{i^*}, h_{i^*})$ for some agent $i^* \in N$. Let $i' \in N$ be the agent of minimum id with $r'_{i'} = \text{rank}_{i'}(y_{i'}, \succ_{i'}, h_{i'}) < r_i$, i.e., $\text{rank}_i(y_i, \succ_i, h_i) = r_i$ for $i = 1, \dots, i' - 1$. This means that the linear program $\text{LP}(\succ; (r_1, \dots, r_{i'-1}, r'_{i'}, p_{i'+1}, \dots, p_n))$ has lottery $\mathbf{y}$ as solution, contradicting the definition of $r_{i'}$ at round i' of the mechanism. $\square$

C Two-sided and stable matchings

In this section we assume that we have two sets of agents N and M , each has n agents. Each agent $i \in N$ (or, $i \in M$) has a quantile parameter $h_i \in [0, 1]$ and a preference order $\succ_i$ over the agents in M (resp. over the agents N). Similar to the one-sided scenario, a matching mechanism takes as input a preference profile $\succ$ and the vector $\mathbf{h}$ of quantile parameters, and returns a lottery over perfect matchings between the sets N and M . We begin by studying a variation of social welfare.

Top-choice maximization Given an instance $(N, M, \succ, \mathbf{h})$ with $h_i = h$ for every $i \in N \cup M$, we design the following polynomial-time algorithm based on maximum weight b -matching (see Algorithm 1). We construct a weighted complete bipartite graph G between the agent sets N and M . Each edge (i, j) such that i is the top choice of agent j and j is the top choice of agent i has weight 2. Each edge (i, j) such that $i \in N$ is the top choice of agent $j \in M$ but agent j is not the top choice of agent i or $i \in N$ is not the top choice of agent $j \in M$ but agent j is the top choice of agent i has weight 1. All other edges have weight 0. Set $b = \lfloor \frac{1}{1-h} \rfloor$. Then, we find a maximum weight b -matching in graph G . The edges of the b -matching will each have probability $1 - h$ and the probabilities of the remaining edges are completed appropriately to form a lottery.

Algorithm 1: Maximize the number of agents that have top-choice as representative

Input: A two-sided matching instance $(N, M, \succ, \mathbf{h})$ with $h_i = h$ for every $i \in N \cup M$

Output: A lottery x that maximizes the number of agents that have their top-choice as representative

Construct a weighted complete bipartite graph $G = (N \cup M, E, wt)$ where

$E = \{\{i, j\} | i \in N, j \in M\}$ and the weight $wt : E \rightarrow \{2, 1, 0\}$ is defined as:

for $\{i, j\} \in E$ **do**

$wt(\{i, j\}) = 2$ if agent i and j are each other's top-choice;
 $wt(\{i, j\}) = 1$, else if agent i is the top-choice of agent j or agent j is the top-choice of i ;
 $wt(\{i, j\}) = 0$, otherwise;

end

Set $b = \lfloor \frac{1}{1-h} \rfloor$;

Let μ = a maximum weight b -matching in G ;

for edge $\{i, j\} \in \mu$ **do**

/* Construct lottery x */
Set $x_{ij} = 1 - h$;

end

Assign probabilities to the remaining edges appropriately such that x is a lottery;

Return x ;

Theorem 19. Given a two-sided matching instance $(N, M, \succ, \mathbf{h})$ with $h_i = h$ for every $i \in N \cup M$, the lottery computed in Algorithm 1 maximizes the number of agents with their top-choice as representative.

Proof. Let x be the lottery returned by the algorithm. Suppose that x does not maximize the number of agents with top-choice as representative. Let y be another lottery that maximizes the number of agents with top-choice as representative. Construct a matching μ' from y as follows. For each agent i , such that its top-choice agent j is its representative in y , add the edge $\{i, j\}$ to μ' . Towards contradiction, we will show that μ' is a b -matching with larger weight than that of μ . Note that $b = \lfloor \frac{1}{1-h} \rfloor$ and each edge that is added to μ' corresponds to a representative. Then, from the definition of representatives, we have that μ' is a b -matching. Now we show that weight of μ' is greater than that of μ . From the construction of the graph G , we have that if both agents in an edge are each other's top-choice, then the edge contributes weight 2 to the total weight and if only one agent in an edge has top-choice as its representative, then it adds weight 1 to the weight of μ' . Thus, for each agent that has its top-choice as representative contributes one to the weight of the b -matching. Therefore, since more agents has top-choice as representative in y , weight of μ' is more than μ , a contradiction. Finally, the algorithm runs in polynomial time as we can construct the graph G and find a weighted b -matching [32] in polynomial time. $\square$

Before we can ask a question similar to Open Problem 2, we ask the more basic question: whether we can extend the above theorem for arbitrary quantile vector $\mathbf{h}$. We could answer it in polynomial time in the one-sided setting.

Open Problem 5. *What is the complexity of maximizing the number of agents with top choice as representative for a given two-sided matching instance with arbitrary quantile vector $\mathbf{h}$?*

In the following, we focus on stability which has been a main area of research in two-sided matching in the last few decades since the seminal paper by Gale and Shapley [22]. We start with the analogous definition of stability in our model.

Definition 20 (Stability). *Given an instance $(N, M, \succ, \mathbf{h})$ of two-sided matching, a lottery x over matchings is stable if for every pair of agents $i \in N$ and $j \in M$, it holds that $\text{rep}(x_i, \succ_i, h_i) \succeq_i j$ or $\text{rep}(x_j, \succ_j, h_j) \succeq_j i$. If this is not the case, i.e., it is $j \succ_i \text{rep}(x_i, \succ_i, h_i)$ and $i \succ_j \text{rep}(x_j, \succ_j, h_j)$, we say that the pair (i, j) is a blocking pair for lottery x .*

We explore the relations between the stability definitions. An *integral matching* is a lottery x where for any pair of agents i and j , it holds that x_{ij} is either 0 or 1. In the next lemma we show that for integral matchings the classical definition of stability is the same as ours.

Lemma 21. *In every two-sided matching instance $(N, M, \succ, \mathbf{h})$, any stable integral matching is also a stable lottery.*

Proof. Let μ be an stable integral matching for the instance $(N, M, \succ)$. We define a lottery x from the matching μ by assigning $x(i, j) = 1$ if $\mu(i) = j$, and assigning zero otherwise, where $i \in N$ and $j \in M$. We show that x is a stable lottery for the instance $(N, M, \succ, \mathbf{h})$. Suppose not and (i, j) is a blocking pair for x . Then, since the preferences are strict ordering, by negation of the condition in Definition 20 we get that there exists a pair $i \in N, j \in M$ such that $j \succeq_i \text{rep}(x_i, \succ_i, h_i)$ and $i \succeq_j \text{rep}(x_j, \succ_j, h_j)$. Observe that $\text{rep}(x_i, \succ_i, h_i) = \mu(i)$ and $\text{rep}(x_j, \succ_j, h_j) = \mu(j)$. Therefore, we get that $j \succeq_i \mu(i)$ and $i \succeq_j \mu(j)$. Then (i, j) is a blocking pair for μ , a contradiction. $\square$

However, a stable lottery captures much more than stable integral matchings.

Lemma 22. *There is a two-sided matching instance $(N, M, \succ, \mathbf{h})$ with $h_i = 1/2$ for every $i \in N \cup M$ such that a stable lottery has an unstable integral matchings in its support.*

Proof. Consider an instance of two sided matching where $N = \{m_1, m_2\}$, $M = \{w_1, w_2\}$, $h_i = 1/2$ for each agent $i \in N \cup M$, and the preference profile is as follows:

$$\begin{array}{ll}
m_1: & w_2 \succ_{m_1} w_1 \\
m_2: & w_2 \succ_{m_2} w_1
\end{array}
\qquad
\begin{array}{ll}
w_1: & m_1 \succ_{w_1} m_2 \\
w_2: & m_2 \succ_{w_2} m_1
\end{array}$$

The lottery x is defined as $x_{m_1 w_2} = 1/3$, $x_{m_1 w_1} = 2/3$, $x_{m_2 w_2} = 2/3$, and $x_{m_2 w_1} = 1/3$. Thus, by symmetry we get, $x_{w_1 m_1} = 2/3$, $x_{w_1 m_2} = 1/3$, $x_{w_2 m_2} = 2/3$, and $x_{w_2 m_1} = 1/3$. Observe that x is stable. However, in a decomposition of the lottery as a convex combination of integral matchings, there must be an integral matching μ that matches m_2 to w_1 and m_1 to w_2 . However, μ is not an integral stable since (m_2, w_2) is blocking. $\square$

Basic Definitions Lotteries for the two sided case can also be viewed as a fractional matching between the two sets N and M . A *fractional matching* $\mu: N \times M \rightarrow \mathcal{R}_{\geq 0}$ is an assignment of non-negative weights to each pair $i \in N$ and $j \in M$ such that $\sum_{j \in M} \mu(i, j) \leq 1$ for each agent $i \in N$ and $\sum_{i \in N} \mu(i, j) \leq 1$ for each agent $j \in M$. Aharoni and Fleiner [4] defined stability for fractional matchings as follows. A pair of agents $i \in N$ and $j \in M$ form an *ordinally blocking pair* in a fractional matching μ if $\sum_{j' \in M: j' \succeq_i j} \mu(i, j') < 1$ and $\sum_{i' \in N: i' \succeq_j i} \mu(i', j) < 1$. If a fractional matching has no ordinally blocking pair, then it is *ordinally stable* (also known as fractionally stable; see [8]).

We observe the following relation between ordinally stable fractional matchings and stable lotteries that follows directly from the definitions.

Observation 23. *Given a two-sided matching instance $(N, M, \succ, \mathbf{h})$ with $h_i = 0$ for every $i \in N \cup M$, a lottery is stable if and only if it is ordinally stable.*

We restate the definition of efficiency which is the same as defined in the one-sided setting.

Definition 24 (Efficiency for lotteries over two-sided matchings). *Given a two-sided instance $(N, M, \succ, \mathbf{h})$, a lottery x over matchings is efficient if there is no other lottery y such that $\text{rep}(y_i, \succ_i, h_i) \succeq_i \text{rep}(x_i, \succ_i, h_i)$ for every $i \in N \cup M$ and there is an agent $i^* \in N \cup M$ such that $\text{rep}(y_{i^*}, \succ_{i^*}, h_{i^*}) \succ_{i^*} \text{rep}(x_{i^*}, \succ_{i^*}, h_{i^*})$.*

Given two lotteries x and y , if representative of every agent in y is at least as good as it is in x and there is an agent whose representative strictly improves in y then y is said to Pareto dominate x . Additionally, y is a *Pareto improvement* over the lottery x .

We show that stability is compatible with efficiency. First we prove the following property.

Lemma 25. *Given a two-sided matching instance $(N, M, \succ, \mathbf{h})$, a stable lottery remains stable after a Pareto improvement.*

Proof. Let y denotes a lottery that is a Pareto improvement over a stable lottery x . That is, where an agent $i^* \in N \cup M$ has strictly better representative in y than in x and the representative of an agent $i' \in N \cup M$ is not worse off in y compared to x . Since x is a stable lottery, we know that for every pair of agents $i \in N$ and $j \in M$, it holds that $\text{rep}(x_i, \succ_i, h_i) \succeq_i j$ or $\text{rep}(x_j, \succ_j, h_j) \succeq_j i$. Without loss of generality suppose $i^* \in N$. Therefore, for every agent $j \in M$, it holds that $\text{rep}(y_{i^*}, \succ_{i^*}, h_{i^*}) \succeq_{i^*} j$. Thus, there is no blocking pair for y containing i^* . The representative of an agent $i' \in N \cup M$ is not worse off in y compared to x . Thus, from stability of x , it holds that for every pair of agents $i \in N \setminus \{i^*\}$ and $j \in M$, it holds that $\text{rep}(y_i, \succ_i, h_i) \succeq_i j$ or $\text{rep}(y_j, \succ_j, h_j) \succeq_j i$. Therefore, y is a stable lottery. $\square$

Finally, we are ready to adapt the PSD mechanism to two-sided matching instances and show the following.

Theorem 26. *Given a two-sided matching instance $(N, M, \succ, \mathbf{h})$, an efficient, stable lottery can be computed in polynomial time.*

Proof. Recall that a stable integral matching is a stable lottery (Lemma 21). To prove the theorem, we employ an algorithm similar to the PSD mechanism in the one-sided scenario starting with a stable integral matching in the two-sided instance.

Given an instance $(N, M, \succ, \mathbf{h})$, we first compute a stable integral matching μ for the instance $(N, M, \succ)$ and compute x as described in Lemma 25. Then solve the following LP to check if there exists a Pareto improvement of x .

For a given two-sided matching instance $(N, M, \succ, \mathbf{h})$ and a vector of rank requirements $\mathbf{r} = (r_1, r_2, \dots, r_n)$ with integer $r_i \in [n]$, the following linear program is feasible if and only if there is a lottery over matchings that satisfy the rank requirement r_i for the representative of each agent $i \in N \cup M$.

$$\begin{aligned} \sum_{g: \text{rank}_i(g) \leq r_i} x_{ig} &\geq 1 - h_i, \quad \forall i \in N \cup M \\ \sum_{g \in M} x_{ig} &= 1, \quad \forall i \in N \\ \sum_{i \in N} x_{ig} &= 1, \quad \forall g \in M \end{aligned}$$

Dropping N, M , and $\mathbf{h}$ from notation, as before, we will refer to this linear program as $\text{LP2}(\succ; \mathbf{r})$. The agents one by one in a fixed order. When the agent $i \in N \cup M$ is considered, the mechanism computes the minimum rank t for agent i so that the linear program $\text{LP2}(\succ; \mathbf{r}_{-i}, t)$ is feasible and updates $\mathbf{r}$ by setting $r_i = t$. After all agents have been considered, an arbitrary lottery over matchings that satisfies $\text{LP2}(\succ; \mathbf{r})$ is returned as the output of the mechanism. First observe that when each agent is considered, the output of LP2 is a Pareto improvement over the previous lottery. Thus by Lemma 25, we maintain the invariant that in each step the lottery is stable. Finally, no agent can improve without making another (previously considered agent) worse-off, thus we reach a Pareto optimal lottery. Clearly, this process runs in polynomial time. $\square$

In the following, we consider tradeoffs between efficiency, stability, and strategyproofness.

Theorem 27. *Given a two-sided matching instance $(N, M, \succ, \mathbf{h})$ with $h_i = 0$ for every $i \in N \cup M$, no stable lottery mechanism can be simultaneously efficient and strategyproof.*

Proof. It is known that an ordinally stable lottery can be written as a convex combination of integral stable matchings and each integrally stable matchings in the support of a ordinally stable lottery is integrally stable and there is a matching μ in the support that is efficient [17]. Thus, using Observation 23, we have that each integral matching in the support of a stable lottery x is integral stable. Next, we show that if x is efficient, then x is integral, i.e., there is one integral stable, efficient matching μ such that $x = \mu$. For every matching μ in the support of a lottery, if there is another matching μ' so that an agent i is worse-off in μ' as compared to μ , then the representative of agent i is the worse mate she has in μ' as $h_i = 0$. Hence, if the two matchings in the support of x are different, then at least one of them Pareto dominates the lottery x and x is not efficient. Thus, any lottery that is efficient is equivalent to an integral matching when $h_i = 0$ for each agent i . Then, the theorem follows by the well-known impossibility result for integral matching mechanisms [37]. $\square$

To overcome the impossibility result and show a positive result for higher values of the quantile parameter, we define a restricted version of lotteries and efficiency for these restricted lotteries. A lottery x has *distinct representatives* if representative of each agent in x is distinct, i.e., for each two distinct agents $i, i' \in N \cup M$ it holds that $\text{rep}(x_i, \succ_i, h_i) \neq \text{rep}(x_{i'}, \succ_{i'}, h_{i'})$.

Observe that a deterministic mechanism produces a matching, i.e., each agent has a distinct representative. Moreover, the representative relation is symmetric. If a lottery has distinct representative, the

former criteria holds, however, the relation is not guaranteed to be symmetric. Thus, this definition brings us closer to a deterministic outcome. Nonetheless, efficient distinct representative outcomes exhibit a contrast to the deterministic case.

Definition 28 (Distinct Representative Efficient Lotteries). *Given a two-sided matching instance $(N, M, \succ, \mathbf{h})$, a lottery x is Distinct Representative Efficient (DR-Efficient) if x has distinct representatives and for each agent there is no other lottery y with distinct representatives such that $\text{rep}(y_i, \succ_i, h_i) \succeq_i \text{rep}(x_i, \succ_i, h_i)$ for every $i \in N$ and there is an agent $i^* \in N$ such that $\text{rep}(y, \succ_i, h_i) \succ_i \text{rep}(x, \succ_i, h_i)$.*

In contrast to stable lotteries or integral stable matchings, we show that a stable and DR-efficient lottery exists. To construct a stable lottery that is DR-efficient we will use the deferred acceptance algorithm, arguably, the most common stable matching algorithm.

Deferred Acceptance Algorithm The *Deferred Acceptance* (DA) algorithm takes as input the preferences of two sets of agents over each other (commonly referred to as men and women). In DA for men, it proceeds in rounds where each round consists of two phases. (i) a proposal phase: Each man without a partner proposes to his most preferred woman from those who have not yet rejected him, and (ii) a rejection phase: each woman receiving multiple proposals rejects all but the one she prefers the most. The algorithm terminates when every agent has a partner. Analogously, one can define DA for women where women are proposing and men accept or reject the proposals.

Gale and Shapley [22] showed that DA always terminates with a stable matching. Moreover, the matching obtained by the men-proposing algorithm is the best for all men among all stable matchings, while being the worst for all women [35].

Thus, we say, for the two set of agents M and N , a DA algorithm for M (resp. N) produces a stable matching that is the *optimal* one for all agents in M (resp. all agents in N) among all stable matchings. Now we are ready to state the following rule that will be used in the next theorem.

We define a mechanism *half-DA* that produces a lottery computed by running deferred acceptance for the set N and for the set M and taking the two stable matchings, each with probability $1/2$.

Theorem 29. *For two-sided matching instances with $h_i \geq 1/2$ for every agent i , the mechanism half-DA is strategyproof and returns a DR-efficient, stable lottery in polynomial time.*

Proof. Let x denote the lottery produced by half-DA. We show that lottery x satisfies the stated properties.

DR-efficiency. First observe that Since $h_i \geq 1/2$ x has distinct representatives and for each agent i , from the definition of x , we have that each agent's representative is their best stable partner. Suppose there exists another lottery y with distinct representatives where an agent i^* gets a better representative than x and no agent $i' \in N \cup M$ is worse off. Without loss of generality, suppose that $i^* \in N$, then we construct a matching μ^* as follows each agent $i \in N$ is matched to its representative in y . Since no agent $i' \in N \cup M$ is worse off than x , then μ^* is stable, contradicting the fact that deferred acceptance for N produces an optimal stable matching for N .

Strategyproofness. It is well known that no agent from the proposing side can obtain a better partner in the outcome deferred acceptance algorithm by misreporting their preference [25]. Observe that each agent in M (and N) receives their optimal partner in M -proposing DA (resp. N -proposing DA). Moreover, in the N -proposing DA, the partner of an agent $j \in M$ cannot be better than j 's partner in M -proposing DA [35]. Thus, in the lottery x , the representative of j remains the same in a lottery produced by our mechanism even after misreporting. This holds true for any agent in $N \cup M$. Thus, no agent has incentive to misreport.

Stability. It is easy to observe that x is stable since the representative of each agent is their best stable integral partner. Thus, we prove the theorem. $\square$

Remark. Although the mechanism in Theorem 29 is DR-efficient, it is not efficient. Consider the following example. Let $N = \{m_1, m_2\}$, $M = \{w_1, w_2\}$, $h_i = 1/2$ for each $i \in N \cup M$, and the preference profile $\succ$ is as follows:

$$\begin{aligned} m_1 : w_1 \succ_{m_1} w_2; & \quad w_1 : m_1 \succ_{w_1} m_2 \\ m_2 : w_1 \succ_{m_2} w_2; & \quad w_2 : m_1 \succ_{w_2} m_2 \end{aligned}$$

Then the lottery x' Pareto dominates the lottery x produced by the half-DA mechanism, where $x'_{ij} = 1/2$ for each $i \in N$ and $j \in M$.

We identified that distinct representative property ensures the existence of an efficiency and strategyproofness mechanism to compute a stable matching. However, we do not know if it is a necessary condition. Thus, we conclude the section with the following question.

Open Problem 6. What is an, as general as possible, sufficient and necessary condition that allows for efficiency, stability, and strategyproofness?