

On Interim Envy-Free Allocation Lotteries

IOANNIS CARAGIANNIS, Aarhus University, Denmark

PANAGIOTIS KANELLOPOULOS, University of Essex, UK

MARIA KYROPOULOU, University of Essex, UK

With very few exceptions, recent research in fair division has mostly focused on deterministic allocations. Deviating from this trend, we study the fairness notion of *interim envy-freeness* (iEF) for lotteries over allocations, which serves as a sweet spot between the too stringent notion of ex-post envy-freeness and the very weak notion of ex-ante envy-freeness. iEF is a natural generalization of envy-freeness to random allocations in the sense that a deterministic envy-free allocation is iEF (when viewed as a degenerate lottery). It is also certainly meaningful as it allows for a richer solution space, which includes solutions that are provably better than envy-freeness according to several criteria. Our analysis relates iEF to other fairness notions as well, and reveals tradeoffs between iEF and efficiency. Even though several of our results apply to general fair division problems, we are particularly interested in instances with equal numbers of agents and items where allocations are perfect matchings of the items to the agents. Envy-freeness can be trivially decided and (when it can be achieved, it) implies full efficiency in this setting. Although computing iEF allocations in matching allocation instances is considerably more challenging, we show how to compute them in polynomial time, while also maximizing several efficiency objectives. Our algorithms use the ellipsoid method for linear programming and efficient solutions to a novel variant of the bipartite matching problem as a separation oracle. We also study the extension of interim envy-freeness notion when payments to or from the agents are allowed. We present a series of results on two optimization problems, including a generalization of the classical rent division problem to random allocations using interim envy-freeness as the solution concept.

CCS Concepts: • **Theory of computation** → **Design and analysis of algorithms**; **Algorithmic game theory and mechanism design**.

Additional Key Words and Phrases: fair division, envy-freeness, random allocations, rent division, price of fairness

ACM Reference Format:

Ioannis Caragiannis, Panagiotis Kanellopoulos, and Maria Kyropoulou. 2021. On Interim Envy-Free Allocation Lotteries. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC '21), July 18–23, 2021, Budapest, Hungary*. ACM, New York, NY, USA, 21 pages. <https://doi.org/10.1145/3465456.3467648>

1 INTRODUCTION

Plenty of situations arise in the real world every day, where assets need to be distributed among individuals. Making sure that everyone gets what they are entitled to is an imperative, yet vague, aspiration that is open to interpretation. Fair division is a research area that deals with problems of distributing assets in a way that is considered fair. Fair allocation problems, that focus on indivisible items, have received considerable attention from the EconCS community recently.

Among the fairness notions that have been proposed to capture the necessity for impartiality and justice, envy-freeness [21, 36] is, without doubt, the prevailing one. Envy-freeness requires

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EC '21, July 18–23, 2021, Budapest, Hungary

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8554-1/21/07...\$15.00

<https://doi.org/10.1145/3465456.3467648>

that each individual, or agent, prefers their own share to anyone else's. However natural and intuitive, though, envy-freeness may not be possible to achieve. In addition, the universality of fair division disputes justifies many different definitions of fairness. Some popular fairness notions in the literature include proportionality and max-min fair share, among others.

The vast majority of the related literature focuses on deterministic allocations. The few recent exceptions (e.g., [3, 7, 22]) that consider random allocations (lotteries or probability distributions over allocations) are either too liberal or too conservative in the fairness concepts they consider. For example, ex-ante envy-freeness compares the random bundle allocated to an agent, in terms of expected valuation, to the random bundle allocated to any other agent. Ex-ante envy-freeness is very weak as a fairness guarantee. Indeed, a lottery that allocates all items to an agent selected uniformly at random is ex-ante envy-free; clearly, such a lottery can hardly be considered fair. On the other extreme, the notion of ex-post envy-freeness requires that every outcome of a random allocation is envy-free. Ex-post envy-freeness is very strict and essentially invalidates the advantages of randomness.

The notion of interim envy-freeness [31] serves as middle ground between ex-post and ex-ante envy-freeness, balancing between the stringency of the constraint and the substance of the fairness guarantee. In particular, consider an instance where a set of indivisible items are to be allocated to a set of agents, and a fixed, publicly known, lottery over allocations. Let an outcome of the lottery be realized and each agent observe only their own allocation. Each agent, then, compares their wealth to the random bundle allocated to any other agent, conditioned upon their own realized allocation. Let, for example, the lottery have many possible outcomes (allocations in its support), but only two of which allocate the bundle a of items to agent 1; let agent 2 obtain bundle b in the first of these outcomes and bundle c in the second one. The interim envy-freeness constraint for agent 1 with respect to agent 2 and bundle a , is satisfied if the value that agent 1 has for a is at least as high as her average value for b and c , according to the probability that agent 2 receives them (the relative probability of the two outcomes in the lottery). If such a constraint is satisfied for every agent, with respect to every other agent and any possible bundle, i.e., everyone's allocated bundle is always worth to them at least what they can estimate anyone else is receiving, then the lottery is said to be interim envy-free.

Interim envy-freeness can be naturally extended to accommodate for payments, similarly to the recent fair allocation literature [24, 27]. It is a known fact that payments can help eliminate envy in the deterministic allocation case, both in the form of subsidies paid to the agents to compensate for an unsatisfactory bundle, and in the form of rent payments paid by the agents to make their allocation look less desirable to others. Rent division [2] is a fundamental fair allocation problem involving payments, where the input consists of a total rent amount, a set of agents, and an equal number of rooms on which the preferences of agents are expressed. The goal is to assign a price to each room so that the room prices sum up to the total rent, and to match agents to rooms so that everyone prefers their own allocation and rent share. Using the interim envy-freeness concept, we consider natural extensions of problems with payments to the random allocation case, both in the rent division and in the subsidy distribution context.

Matching allocation instances, as in the rent division setting just discussed, are relevant in many applications; hence, we partially focus on this case. An important technical advantage is that such instances allow for an easy computation of (deterministic) envy-free allocations, as opposed to general allocation instances, for which relevant problems are typically NP-hard. However, allowing randomization seems to make the situation considerably more complex. Interestingly, as we will see, the added complication still allows for positive computational results related to interim envy-free lotteries.

1.1 Overview and Significance of Our Contribution

To the best of our knowledge, interim envy-freeness (iEF) has not received any attention by the EconCS community. We justify its importance as a fairness notion for lotteries of allocations, by demonstrating a rich menu of interesting properties it enjoys. First, we relate it to the most important fairness properties for deterministic allocations and lotteries. In terms of strength as a fairness property, iEF is proved to lie between proportionality and envy-freeness in the following way. Clearly, when viewed as a degenerate lottery, any envy-free allocation is iEF. Also, every iEF lottery is defined over proportional allocations. These implications are shown to be strict in a strong sense. We show that there are allocation instances that admit proportional allocations but no iEF lottery, and instances that admit iEF lotteries but no envy-free allocation. Compared to fairness properties for lotteries, iEF lies between ex-ante envy-freeness (which can be always attained trivially) and ex-post envy-freeness (which is too restrictive). These findings and observations appear in Section 3.

Our next goal is to explore the trade-offs between iEF and economic efficiency (Section 4). We pay special attention to matching allocation instances where envy-freeness implies Pareto-efficiency and optimal utilitarian, egalitarian, and average Nash social welfare. A careful interpretation of these facts reveals that envy-freeness is a very restrictive fairness property. In contrast, as less restrictive, iEF lotteries may not be Pareto-efficient and can furthermore produce allocations of suboptimal social welfare. We provide tight or almost tight bounds on the price of iEF with respect to the utilitarian, egalitarian, and average Nash social welfare. These bounds suggest that iEF allocations can have social welfare that is up to $\Theta(n)$ times far from optimal, where n is the number of agents.

Bounds on the price of iEF give only rough estimates of the best social welfare of iEF lotteries. We present polynomial-time algorithms for computing iEF lotteries that maximize the utilitarian, egalitarian, and log-Nash social welfare. Our algorithms follow a general template that can be briefly described as follows. The problem of computing an iEF lottery of maximum social welfare is formulated as a linear program. This linear program has exponentially many variables; to solve it, we exploit the execution of the ellipsoid method to its dual. As the dual linear program has exponentially many constraints, the ellipsoid method needs access to polynomial-time separation oracles that check whether the dual variables violate the dual constraints or not. We design such separation oracles by exploiting connections to maximum edge-pair-weighted bipartite perfect matching (2EBM), a novel (to the best of our knowledge) combinatorial optimization problem that involves perfect matchings in bipartite graphs. We show how to solve 2EBM in polynomial time by exploiting an elegant lemma by Cruse [19] on decompositions of doubly-stochastic centrosymmetric matrices. We believe that 2EBM is a natural combinatorial optimization problem of independent interest and with applications in other contexts. These computational results appear in Section 5 and constitute the most technically intriguing results in the paper.

Finally, we extend the definition of interim envy-freeness to accommodate for settings where monetary transfers (payments) are allowed. We define and study two related optimization problems. In subsidy minimization, which is motivated by a similar problem for deterministic allocations that was studied recently, we seek iEF pairs of lotteries and payments to the agents so that the total expected amount of payments is minimized. In utility maximization, which extends the well-known rent division problem, we seek iEF pairs of lotteries and payments that are collected from the agents and contribute to a fixed rent; the objective is to maximize the minimum expected agent utility. We consider different types of payments depending on whether the payments are agent-specific, bundle-specific, or unconstrained (i.e., specific to agents and allocations). iEF is proved to be considerably more powerful than envy-freeness, allowing for much better solutions to the two optimization

problems compared to their deterministic counterparts. We also showcase the importance of both agent-specific and bundle-specific payments by showing that they are incomparable to each other, in the context of the two optimization problems. By applying our computational template, we present efficient algorithms that compute optimal solutions to subsidy minimization and utility maximization using unconstrained payments, violating the iEF condition only marginally. These results are presented in Section 6. We believe that they attest to the significance of interim envy-freeness too and will motivate further study.

Due to lack of space, several proofs and intermediate statements have been omitted. They will appear in the final version of the paper.

1.2 Further Related Work

Previous work on randomness in allocation problems is clearly related to ours. Aziz [6] discusses the benefits of randomization in social choice settings, including fair division, thus supporting relevant studies despite the related challenges. Gajdos and Tallon [23] study *ex-ante* and *ex-post* notions of fairness in a setting with an inherent uncertainty imposed by the environment. More recently, Aleksandrov et al. [3] analyze the *ex-post* and *ex-ante* envy-freeness guarantees of particular algorithms for online fair division. Freeman et al. [22] study the possibility of achieving *ex-ante* and *ex-post* fairness guarantees simultaneously in the classical fair allocation setting. For example, they show that there always exists an *ex-ante* envy-free lottery, with all allocations in its support satisfying a relaxed fairness property known as envy-freeness up to one item (EF1; see [13]). They furthermore show that such lotteries can be computed efficiently. In very recent work, Aziz [7] strengthens some of these results.

Interim envy-freeness was first defined by Palfrey and Srivastava [31] (see also [20, 37]). Even though the intuition behind the iEF notion in those papers coincides with ours, their settings are different. They are more complex in the sense that the value of an agent for a bundle may depend on the allocation, but they are more restrictive as the lottery probabilities are fixed in advance. Our setting is more suitable to formulate and study existence and computational questions.

Randomness is an important design tool for mechanisms that compute solutions in matching allocation instances. In a slightly different context than ours (e.g., see [28]), agents are assumed to have ordinal preferences on the items. Mechanisms such as the probabilistic serial rule [11] introduce randomness to avoid the discrimination between agents, and achieve *ex-ante* fairness guarantees for all cardinal valuations that are compatible to the ordinal preferences. Other rules in this line of research include random priority [1], vigilant eating [8], and several extensions of the probabilistic serial rule [14].

Previous work on subsidy minimization has focused on envy-freeness as the solution concept. The objective is to compute an allocation and appropriate payments to the agents so that agents are non-*envious* for the combinations of payments and bundles of items they receive. The work of Maskin [30] seems to be the first treatment of the problem, without optimizing subsidies though. The notion of *envy-freeability* refers to an allocation that can become *envy-free* when paired with appropriate payments. Halpern and Shah [27] present a characterization that indicates that *envy-freeability* is strongly connected to a *no-positive-cycle* property in an appropriately defined *envy graph*. Alternatively, *envy-freeable* allocations maximize the utilitarian social welfare with respect to all bundle redistributions among the agents. Among other results, Halpern and Shah [27] aim to bound the amount of subsidies assuming that all agent valuations are in $[0, 1]$. They conjecture that subsidies of $n - 1$ suffice; an even stronger version of the conjecture is proved by Brustle et al. [12]. Complexity results for subsidy minimization are presented by Caragiannis and Ioannidis [15]. Another study that blends fairness (including *envy-freeness*) with payments is the work of Chevalere et al. [18] on distributed allocations of items.

Rent division had attracted attention much before subsidy minimization [5, 33, 34]. A very similar characterization of envy-freeable allocations [4, 35] like the one mentioned above, has allowed for a simple solution in [24] to the problems of maximizing the minimum agent utility and minimizing the disparity between agent utilities. Earlier, the papers [2, 26, 29] present algorithms for computing envy-free rent divisions without considering optimization criteria.

2 PRELIMINARIES

An instance of an allocation problem consists of a set \mathcal{N} of n agents and a set \mathcal{I} of m items. Agent $i \in \mathcal{N}$ has valuation $v_i(j)$ for item $j \in \mathcal{I}$. By abusing notation, we use $v_i(S)$ to denote the valuation of agent i for the set (or bundle) of items S . We assume that valuations are non-negative and additive, i.e., $v_i(S) = \sum_{j \in S} v_i(j)$. We remark, though, that several of our results (including the concept of interim envy-freeness, which we define later in this section) carry over to more general valuations. Furthermore, even though this is rarely required for our positive statements, in our examples we use normalized valuations satisfying $\sum_{j \in \mathcal{I}} v_i(j) = 1$ for every agent $i \in \mathcal{N}$.

An allocation $A = (A_1, A_2, \dots, A_n)$ of the items in \mathcal{I} to the agents of \mathcal{N} is simply a partition of the items of \mathcal{I} into n bundles A_1, A_2, \dots, A_n , with the understanding that agent $i \in \mathcal{N}$ gets the bundle of items A_i . An allocation $A = (A_1, A_2, \dots, A_n)$ is *envy-free* (EF) if $v_i(A_i) \geq v_i(A_k)$ for every pair of agents i and k . In words, the allocation A is envy-free if no agent prefers the bundle of items that has been allocated to some other agent to her own. The allocation A is *proportional* if $v_i(A_i) \geq \frac{1}{n} v_i(\mathcal{I})$. An instance may not admit any envy-free or proportional allocation; to see why, consider an instance in which all agents have a positive valuation of 1 for a single item (and zero value for any other item). It is well-known that, due to additivity, an envy-free allocation is always proportional.

Envy-freeness is defined accordingly if monetary transfers are allowed. In particular, a pair of an allocation A and a vector \mathbf{p} consisting of a payment p_i to each agent $i \in \mathcal{N}$ is envy-free with payments if $v_i(A_i) + p_i \geq v_i(A_k) + p_k$ for every pair of agents i and k . The term *envy-freeable* refers to an allocation that can become envy-free with an appropriate payment vector. Depending on the setting, payments can be restricted to be non-negative (e.g., representing subsidies that are given to the agents [12, 15, 27]) or non-positive (e.g., when payments are collected from the agents, like in the rent division problem [24]).

In addition to their fairness properties, allocations are typically assessed in terms of their efficiency. We say that an allocation $A = (A_1, A_2, \dots, A_n)$ is Pareto-efficient if there is no other allocation $A' = (A'_1, A'_2, \dots, A'_n)$ with $v_i(A'_i) \geq v_i(A_i)$ for every agent $i \in \mathcal{N}$, with the inequality being strict for at least one agent of \mathcal{N} . The term social welfare is typically used to assign a cardinal score that characterizes the efficiency of an allocation. Among the several social welfare notions, the utilitarian, egalitarian, and Nash social welfare are the three most prominent. We use the notation $U(A)$, $E(A)$, $\text{avN}(A)$, and $\text{lgN}(A)$ to refer to the utilitarian, egalitarian, average Nash, and log-Nash social welfare, respectively, of an allocation $A = (A_1, \dots, A_n)$; the corresponding efficiency scores are defined as follows:

$$U(A) = \sum_{i \in \mathcal{N}} v_i(A_i), \quad E(A) = \min_{i \in \mathcal{N}} v_i(A_i), \quad \text{avN}(A) = \left(\prod_{i \in \mathcal{N}} v_i(A_i) \right)^{1/n}, \quad \text{lgN}(A) = \sum_{i \in \mathcal{N}} \ln v_i(A_i).$$

The *price of fairness*, introduced independently in [9] and [16], refers to a class of notions that aim to quantify trade-offs between fairness and efficiency. For example, the price of envy-freeness with respect to the utilitarian social welfare for an allocation instance (that admits at least one envy-free allocation) is the ratio of the optimal utilitarian social welfare of the instance over the

utilitarian social welfare of the best envy-free allocation. Different price of fairness notions follow by selecting different fairness concepts and social welfare definitions.

We are particularly interested in *matching allocation instances*, in which the number of agents is equal to the number of items. Our assumptions (e.g., for non-negative valuations) imply that the only reasonably fair (e.g., proportional) allocations should then assign (or match) exactly one item to each agent. We refer to such allocations as *matchings*.¹ Notice that, an envy-free matching must allocate to each agent her most-valued item. As such, whenever an envy-free allocation exists in a matching instance, it is Pareto-efficient and maximizes the social welfare, according to all definitions of social welfare mentioned above. Hence, the price of envy-freeness is trivially 1 in this case (with respect to all the social welfare definitions given above).

Random allocations and interim envy-freeness

We consider random allocations that are produced according to *lotteries* (or probability distributions). The lottery \mathbf{Q} over allocations of the items of \mathcal{I} to the agents of \mathcal{N} is ex-ante envy-free if $\mathbb{E}_{A \sim \mathbf{Q}}[v_i(A_i)] \geq \mathbb{E}_{A \sim \mathbf{Q}}[v_i(A_k)]$ for every pair of agents $i, k \in \mathcal{N}$. \mathbf{Q} is ex-post envy-free if any allocation it produces with positive probability is envy-free (or, in other words, if all allocations in the *support* of \mathbf{Q} are envy-free).

We now provide the formal definition of the central concept of this paper. We say that a lottery \mathbf{Q} over allocations is *interim envy-free* (iEF) if for any pair of agents $i, k \in \mathcal{N}$ and any possible bundle of items S that agent i can get in a random allocation produced by \mathbf{Q} , it holds

$$v_i(S) \geq \mathbb{E}_{A \sim \mathbf{Q}}[v_i(A_k) | A_i = S].$$

The definition of iEF requires that the value agent i has when she gets a bundle S is at least as high as the average value she has for the bundle that agent k gets, conditioned on i 's allocation.

We extend the notion of interim envy-freeness to pairs of lotteries over allocations and payments to/from the agents, in an analogous way that recent work has defined envy-freeness with payments. In fact, we differentiate between different payment schemes with respect to whether payments are per agent (A-payments), per bundle (B-payments), or per allocation and agent (C-payments). We similarly extend the notion of price of fairness and envy-freeability to the case of iEF. We postpone providing formal definitions for the corresponding sections that these notions are being examined.

3 INTERIM ENVY-FREENESS VS. OTHER FAIRNESS NOTIONS

In this section we compare interim envy-freeness with other fairness notions, with the aim to identify possible fairness implications. The first implication follows easily by the definitions and has been observed before in more general contexts than ours (e.g., see [37]).

LEMMA 3.1. *Any iEF lottery \mathbf{Q} is ex-ante envy-free.*

PROOF. Indeed, using the definitions of iEF, ex-ante envy-freeness and well-known properties of random variables, we have

$$\begin{aligned} \mathbb{E}_{A \sim \mathbf{Q}}[v_i(A_i)] &= \sum_{S \subseteq \mathcal{I}} v_i(S) \cdot \Pr_{A \sim \mathbf{Q}}[A_i = S] \geq \sum_{S \subseteq \mathcal{I}} \mathbb{E}_{A \sim \mathbf{Q}}[v_i(A_k) | A_i = S] \cdot \Pr_{A \sim \mathbf{Q}}[A_i = S] \\ &= \mathbb{E}_{A \sim \mathbf{Q}}[v_i(A_k)], \end{aligned}$$

for every pair of agents i and k . □

¹We typically use the small letter b to denote a matching, instead of the usual notation of A for allocations in general instances.

An even simpler observation is that any lottery that deterministically produces an envy-free allocation A is trivially iEF. Indeed, A_i is the only bundle that can be given to agent i , who weakly prefers it to the bundle A_k that is allocated to agent k . Trivially, $\mathbb{E}_{A \sim Q}[v_i(A_k)|A_i = S] = v_i(A_k)$ and the iEF condition is identical to the envy-freeness condition $v_i(A_i) \geq v_i(A_k)$. We can slightly extend this argument to obtain the following implication.

LEMMA 3.2. *Any ex-post envy-free lottery Q is iEF.*

However, the opposite is not true; we show below that the existence of an iEF allocation does not imply the existence of an EF allocation. This is important as it indicates that the set of iEF allocations is larger than those of EF ones.

LEMMA 3.3. *There exist allocation instances with an iEF lottery but with no EF allocation.*

PROOF. Consider the matching allocation instance at the left of Table 1 and the lottery Q which returns matchings a - b - c and a - c - b , with probability $1/2$ each. As agents 2 and 3 have maximum possible value in each of the two matchings produced by Q , EF and, consequently, iEF conditions for them are satisfied. To see that the iEF condition is satisfied for agent 1, observe that she is allocated item a in both matchings in the support of Q , for which she has a value of $1/3$. Agent 2 (and, similarly, agent 3) gets item b with probability $1/2$ and item c with probability $1/2$. Agent 1's average value for the item agent 2 (or agent 3) gets is $2/3 \cdot 1/2 + 0 \cdot 1/2 = 1/3$. Hence, the iEF condition for agent 1 with respect to agent 2 (and, similarly, for agent 1 with respect to 3) is satisfied. The proof that the lottery Q is iEF is complete.

	a	b	c
1	$1/3$	$2/3$	0
2	0	$1/2$	$1/2$
3	0	$1/2$	$1/2$

	a	b	c
1	$1/3$	$2/3$	0
2	0	$2/3$	$1/3$
3	$1/4$	$1/2$	$1/4$

Table 1. The two matching instances that are used in the proofs of Lemmas 3.3 and 3.5 to distinguish between iEF, EF, and proportionality. Throughout the paper, we consider several examples with three agents and items a , b , and c . A concise notation like a - c - b is used to represent the matching in which agents 1, 2, and 3 get items a , c , and b , respectively.

In the same example, it can be easily seen that there is no EF allocation. Indeed, as the only agent who has positive value for item a is agent 1, agent 1 should get this item and be envious of the agent who gets item b . The proof of the lemma is complete. \square

Our next lemma relates iEF to proportionality and is used extensively in our proofs.

LEMMA 3.4. *Any allocation in the support of an iEF lottery is proportional.*

PROOF. Consider an iEF lottery Q and any agent $i \in \mathcal{N}$. By the definition of iEF, we have that for any allocation in the support of Q where agent i gets the bundle of items S , it holds that $v_i(S) \geq \mathbb{E}_{A \sim Q}[v_i(A_k)|A_i = S]$ for each other agent k . By summing up over all other agents we get

$$(n-1)v_i(S) \geq \sum_{k \neq i} \mathbb{E}_{A \sim Q}[v_i(A_k)|A_i = S] = \mathbb{E}_{A \sim Q}[\sum_{k \neq i} v_i(A_k)|A_i = S] = v_i(\mathcal{I} \setminus S).$$

By adding $v_i(S)$ to both sides of the above inequality and rearranging, we get $v_i(S) \geq \frac{1}{n}v_i(\mathcal{I})$, implying that any allocation in the support of Q is proportional. \square

However, iEF is a stronger property than proportionality as the next lemma shows.

LEMMA 3.5. *There exist allocation instances with a proportional allocation but with no iEF lottery.*

PROOF. Consider the instance at the right of Table 1. In this instance, allocation a - c - b is the only proportional allocation. Hence, by Lemma 3.4, to show that no iEF lottery exists, it suffices to consider only the (lottery that deterministically returns) allocation a - c - b . In this allocation, agents 1 and 2 are envious of agent 3, contradicting the iEF requirement. \square

Our next lemma quantifies the disparity between envy-freeness and interim envy-freeness; the proof exploits Lemma 3.5.

LEMMA 3.6. *The maximum envy at any allocation in the support of an iEF lottery, when the agent valuations are normalized, can be as high as $1 - \frac{2}{n}$ and this is tight.*

4 INTERIM ENVY-FREENESS VS. EFFICIENCY

We now explore tradeoffs between interim envy-freeness and efficiency. Two well-studied refinements of Pareto-efficiency are relevant for lotteries of allocations: ex-ante and ex-post Pareto-efficiency. A lottery Q over allocations is ex-ante Pareto-efficient if there exists no other lottery Q' such that $\mathbb{E}_{A \sim Q'}[v_i(A_i)] \geq \mathbb{E}_{A \sim Q}[v_i(A_i)]$ for every agent $i \in \mathcal{N}$, with the inequality being strict for at least one agent of \mathcal{N} . A lottery is ex-post Pareto-efficient if all allocations in its support are Pareto-efficient. It is well-known that ex-ante Pareto-efficiency implies ex-post Pareto-efficiency.

For allocation instances with two agents, the allocations in the support of an iEF lottery are envy-free and, thus (as observed in Section 2), Pareto-efficient. This is due to the fact that any allocation in the support of an iEF lottery is proportional (by Lemma 3.4) and hence envy-free, since there are only two agents. This implies that an iEF lottery is ex-post and ex-ante Pareto-efficient.² In the following, we show that this may not be the case in instances with more agents.

THEOREM 4.1. *There exist matching instances with $n \geq 3$ agents in which no iEF lottery is ex-post (and, consequently, ex-ante) Pareto-efficient.*

To assess the impact of fairness in random allocations to social welfare, we need to extend the price of fairness definition to lotteries. We do so implicitly here by defining the price of iEF (one can similarly define, e.g., the price of ex-ante envy-freeness). We say that the price of iEF with respect to a social welfare measure is the worst-case ratio over all allocation instances with at least one iEF lottery, of the optimal social welfare in the instance over the expected social welfare of the best iEF lottery (where “best” is defined with respect to this social welfare measure).

In our next theorems, we bound the price of iEF with respect to different social welfare notions. In the proof of our upper bounds, we consider normalized valuations. This is a typical assumption in the related literature as well, e.g., see [16]. This assumption is not necessary for average Nash social welfare.

THEOREM 4.2. *The price of iEF with respect to the utilitarian, egalitarian, and average Nash social welfare is at most n , when the agent valuations are normalized.*

PROOF. The proof follows by Lemma 3.4, which implies that the valuation of each agent in any allocation in the support of an iEF lottery is at least $1/n$. Then, the utilitarian, egalitarian, and average Nash social welfare is at least 1, $1/n$, and $1/n$, respectively, while the corresponding optimal values are at most n , 1, and 1, respectively. \square

The next two statements indicate that our price of iEF upper bounds with respect to utilitarian and egalitarian social welfare are asymptotically tight.

²It can be easily seen that ex-post and ex-ante Pareto-efficiency coincide for matching allocation instances with two agents. This is not true in general.

THEOREM 4.3. *The price of iEF with respect to the utilitarian social welfare is at least $\Omega(n)$.*

PROOF. Let $\epsilon > 0$ be negligibly small and $k \geq 2$ be an integer. We use the following matching instance with $n = 2k$ agents and items. For $i = 1, 2, \dots, k$, agent i has value $\frac{k}{k+1}$ for item i , value $\frac{1}{k+1}$ for item $i + k$, and value 0 for any other item. For $i = k + 1, k + 2, \dots, 2k$, agent i has value $\frac{1}{2k} + \epsilon$ for items $1, 2, \dots, k$ and value $\frac{1}{2k} - \epsilon$ for items $k + 1, \dots, 2k$.

An optimal allocation has utilitarian social welfare (at least) $\frac{k^2}{k+1} + \frac{1}{2} - k\epsilon$. To see why, consider the allocation in which agent i gets item i for $i = 1, 2, \dots, 2k$. We now claim that no iEF lottery \mathbf{Q} over allocations has welfare higher than $\frac{k}{k+1} + \frac{1}{2} + k\epsilon$. The lower bound on the price of iEF will follow by the relation between n and k (and by taking ϵ to be sufficiently small). Indeed, by Lemma 3.4, the support of \mathbf{Q} should consist of allocations in which agents $k + 1, k + 2, \dots, 2k$ get items $1, 2, \dots, k$ for a total value of $\frac{1}{2} + k\epsilon$. Then, the maximum value each of the agents $1, 2, \dots, k$ gets from the items $k + 1, \dots, 2k$ is $\frac{1}{k+1}$.

It remains to present such a lottery \mathbf{Q} . It suffices to assign item $i + k$ to agent i for $i = 1, 2, \dots, k$ and assign uniformly at random the items $1, 2, \dots, k$ to the agents $k + 1, \dots, 2k$. Clearly, agents $k + 1, \dots, 2k$ are not envious. For $i = 1, \dots, k$, agent i has value $\frac{1}{k+1}$. Her expected value for the item of another agent ℓ is 0 if ℓ is one of the k first agents besides i and $\frac{1}{k+1}$ if ℓ is one of the k last agents. Notice that in the latter case, agent ℓ gets item i (for which agent i has value $\frac{k}{k+1}$) with probability $1/k$, while she gets items for which agent i has no value otherwise. \square

THEOREM 4.4. *The price of iEF with respect to the egalitarian social welfare is at least $\Omega(n)$.*

A very similar proof to the one of Theorem 4.3 yields our best (albeit not known to be tight) lower bound for the price of iEF with respect to the average Nash social welfare.

THEOREM 4.5. *The price of iEF with respect to the average Nash social welfare is at least $\Omega(\sqrt{n})$.*

5 COMPUTING EFFICIENT INTERIM ENVY-FREE ALLOCATIONS

We devote this section to proving Theorem 5.1 and show how we can compute an iEF lottery of maximum expected social welfare efficiently in the case of matching instances. We remark that, in the case of non-matching instances (where the number of items is larger than the number of agents), deciding whether an iEF lottery exists is an NP-hard problem.³

THEOREM 5.1. *For matching instances, an iEF lottery of maximum expected utilitarian, egalitarian, or log-Nash social welfare (if one exists) can be computed in polynomial time in terms of the number of agents.*

Our algorithms use linear programming. Let \mathcal{M} be the set of all possible perfect matchings between the agents in \mathcal{N} and the items in \mathcal{I} (more formally, \mathcal{M} is the set of all perfect matchings in the complete bipartite graph $G = (\mathcal{N}, \mathcal{I}, \mathcal{N} \times \mathcal{I})$). For agent $i \in \mathcal{N}$ and item $j \in \mathcal{I}$, denote by \mathcal{M}_{ij} the set of matchings from \mathcal{M} in which item j is assigned to agent i . Also, for a matching $b \in \mathcal{M}$ and an agent $k \in \mathcal{N}$, $b(k)$ denotes the item of \mathcal{I} to which agent k is matched in b . Then, an iEF

³To see this, consider the case of two agents and recall that (i) an iEF lottery is a lottery over proportional allocations (Lemma 3.4), and (ii) a proportional allocation is envy-free and (trivially) iEF. Hence, the existence of an iEF lottery is equivalent to the existence of a proportional/EF allocation in the case of two agents. Now, if we consider instances with identical valuations (in which both agents have value $v(g)$ for item g), deciding whether an iEF lottery exists is equivalent to deciding PARTITION, a well-known NP-hard problem.

lottery can be computed as the solution to the following linear program.

$$\begin{aligned}
 & \text{maximize} && \sum_{b \in \mathcal{M}} x(b) \cdot \text{SW}(b) \\
 & \text{subject to} && \sum_{b \in \mathcal{M}_{ij}} x(b) \cdot (v_i(j) - v_i(b(k))) \geq 0, \quad i \in \mathcal{N}, j \in \mathcal{I}, k \in \mathcal{N} \setminus \{i\} \\
 & && \sum_{b \in \mathcal{M}} x(b) = 1 \\
 & && x(b) \geq 0, \quad b \in \mathcal{M}
 \end{aligned} \tag{1}$$

The variables of the linear program are the probabilities $x(b)$, for every matching $b \in \mathcal{M}$, with which the lottery produces matching b . Together with the non-negativity constraints on \mathbf{x} , the second constraint $\sum_{b \in \mathcal{M}} x(b) = 1$ requires that the vector of probabilities $\mathbf{x} = (x(b))_{b \in \mathcal{M}}$ defines a lottery over all matchings of \mathcal{M} . The notation $\text{SW}(b)$ is used here to refer generally to the social welfare of matching b . We will specifically replace SW by U , E , and lgN later. The objective of the linear program is to maximize the expected social welfare $\mathbb{E}_{b \sim \mathbf{x}}[\text{SW}(b)]$ or, equivalently, the quantity $\sum_{b \in \mathcal{M}} x(b) \cdot \text{SW}(b)$.

The first set of constraints represent the iEF conditions. Indeed, the constraint is clearly true for every agent $i \in \mathcal{N}$ and item $j \in \mathcal{I}$ that is never assigned to agent i under \mathbf{x} (i.e., when $\Pr_{b \sim \mathbf{x}}[b(i) = j] = 0$). We will also show that this is the case when $\Pr_{b \sim \mathbf{x}}[b(i) = j] > 0$ as well. For agent $k \in \mathcal{N} \setminus \{i\}$, interim envy-freeness requires that

$$v_i(j) \geq \mathbb{E}_{b \sim \mathbf{x}}[v_i(b(k)) | b(i) = j]. \tag{2}$$

By multiplying the left-hand-side of (2) with $\Pr_{b \sim \mathbf{x}}[b(i) = j]$, we get

$$v_i(j) \cdot \Pr_{b \sim \mathbf{x}}[b(i) = j] = \sum_{b \in \mathcal{M}_{ij}} x(b) \cdot v_i(j),$$

and by doing the same with the right-hand-side, we have

$$\mathbb{E}_{b \sim \mathbf{x}}[v_i(b(k)) | b(i) = j] \cdot \Pr_{b \sim \mathbf{x}}[b(i) = j] = \sum_{b \in \mathcal{M}_{ij}} x(b) \cdot v_i(b(k)).$$

Hence, inequality (2) is equivalent to the first constraint of the linear program (1).

The linear program (1) has exponentially many variables, i.e., one variable for each of the $n!$ different matchings of \mathcal{M} . To solve it efficiently, we will resort to its dual linear program

$$\begin{aligned}
 & \text{maximize} && z \\
 & \text{subject to} && z + \sum_{\substack{(i,j) \in \mathcal{b} \\ k \in \mathcal{N} \setminus \{i\}}} (v_i(j) - v_i(b(k))) \cdot y(i, j, k) + \text{SW}(b) \leq 0, \quad b \in \mathcal{M} \\
 & && y(i, j, k) \geq 0, \quad i \in \mathcal{N}, j \in \mathcal{I}, k \in \mathcal{N} \setminus \{i\}
 \end{aligned} \tag{3}$$

The dual linear program (3) has polynomially many variables and exponentially many constraints. Fortunately, we will be able to solve it using the ellipsoid method [25, 32]. To do so, all we need is a polynomial-time separation oracle, which takes as input values for the dual variables z and $y(i, j, k)$ for all triplets (i, j, k) consisting of agent $i \in \mathcal{N}$, item $j \in \mathcal{I}$, and agent $k \in \mathcal{N} \setminus \{i\}$, and either computes a matching b^* for which a particular constraint is violated, or correctly concludes that no constraint of the dual linear program (3) is violated.

Let us briefly remind the reader how solving the dual linear program using the ellipsoid method can give us an efficient solution to the primal linear program as well; a more detailed discussion can be found in [25, 32]. To solve the dual linear program, the ellipsoid method will make only polynomially many calls to the separation oracle. This is due to the fact that, among the exponentially many constraints, the ones that really constrain the variables of the dual linear program are very

few; the rest are just redundant. Then, after having kept track of the execution of the ellipsoid method on the dual linear program, the primal linear program can be simplified by setting implicitly to 0 all variables that correspond to dual constraints that were not returned as violated ones by the calls of the separation oracle during the execution of the ellipsoid method. As a final step, the solution of the simplified primal linear program (which is now of polynomial size) will give us the solution \mathbf{x} ; this will have only polynomially-many matchings in its support.

In the rest of this section, we will show how to design such separation oracles for the dual linear program (3) when we use the utilitarian, egalitarian, or log-Nash definition of the social welfare. Our separation oracles essentially solve instances of a novel variation of the maximum bipartite matching problem. We believe that this can be of independent interest, with applications in many different contexts.

The maximum edge-pair-weighted perfect bipartite matching

Instances of the maximum edge-pair-weighted perfect bipartite matching problem (or, 2EBM, for short) consist of the complete bipartite graph $G = (\mathcal{N}, \mathcal{I}, \mathcal{N} \times \mathcal{I})$ with a weighting function ψ that assigns weight $\psi(e_1, e_2)$ to every ordered pair of non-incident edges e_1 and e_2 from $\mathcal{N} \times \mathcal{I}$. The objective is to compute a perfect matching $b \in \mathcal{M}$ so that the total weight over all edge-pairs of b , denoted by

$$\Psi(b) = \sum_{(i,j) \in b} \sum_{\substack{(k,\ell) \in b: \\ k \neq i}} \psi((i,j), (k,\ell)), \quad (4)$$

is maximized.⁴

Let \mathcal{X} be the set of quadruples (i, j, k, ℓ) where $i, k \in \mathcal{N}$ and $j, \ell \in \mathcal{I}$, with $i \neq k$ and $j \neq \ell$. We can view such a quadruple as the ordered pair of edges (i, j) and (k, ℓ) in the input graph G . Essentially, the quadruples of \mathcal{X} correspond to all possible ordered pairs of different edges in the input graph. To compute a perfect matching $b \in \mathcal{M}$ of maximum total edge-pair weight, we will use the following integer linear program with $\Theta(n^4)$ variables and constraints (where n is the number of agents and items). We remark that, from now on, we simplify the notation $\psi((i, j), (k, \ell))$ and use $\psi(i, j, k, \ell)$ instead.

$$\begin{aligned} & \text{maximize} && \sum_{(i,j,k,\ell) \in \mathcal{X}} t(i, j, k, \ell) \cdot \psi(i, j, k, \ell) \\ & \text{subject to} && \sum_{\substack{j,\ell \in \mathcal{I}: \\ (i,j,k,\ell) \in \mathcal{X}}} t(i, j, k, \ell) = 1, && i \in \mathcal{N}, k \in \mathcal{N} \setminus \{i\} \\ & && \sum_{\substack{i,k \in \mathcal{N}: \\ (i,j,k,\ell) \in \mathcal{X}}} t(i, j, k, \ell) = 1, && j \in \mathcal{I}, \ell \in \mathcal{I} \setminus \{j\} \\ & && t(i, j, k, \ell) = t(k, \ell, i, j), && (i, j, k, \ell) \in \mathcal{X} \\ & && t(i, j, k, \ell) \in \{0, 1\}, && (i, j, k, \ell) \in \mathcal{X} \end{aligned} \quad (5)$$

For a quadruple $(i, j, k, \ell) \in \mathcal{X}$, the variable $t(i, j, k, \ell)$ indicates whether both edges (i, j) and (k, ℓ) belong to the perfect matching ($t(i, j, k, \ell) = 1$) or not ($t(i, j, k, \ell) = 0$). The first constraint indicates that among all edge pairs with endpoints at agent nodes i and k , exactly one has both its edges in the matching. Similarly, the second constraint indicates that among all edge pairs with endpoints at item nodes j and ℓ , exactly one has both its edges in the matching. The third constraint ensures symmetry of the variables so that they are consistent to our interpretation.

⁴We remark that the problem of computing a perfect matching of maximum total edge weight in an edge-weighted complete bipartite graph with edge weight $w(e)$ for each edge e , is equivalent to 2EBM by defining the edge-pair weights of the latter as $\psi(e, e') = \frac{w(e)}{n-1}$ for every ordered pair of edges e and e' .

We relax the integrality constraint of (5) and replace it by

$$t(i, j, k, \ell) \geq 0, \quad (i, j, k, \ell) \in \mathcal{X}. \quad (6)$$

Then, we compute an extreme solution of the resulting linear program (e.g., again, using the ellipsoid method [25, 32]). We claim (in Lemma 5.4) that this solution is integral, i.e., all variables have values either 0 or 1, and are, hence, solutions to the integer linear program (5) and, consequently, to our maximum edge-pair-weighted perfect bipartite matching problem.

To prove this, we can view the solution \mathbf{t} of the relaxation of the linear program (5) as a square matrix T . In this matrix, each row corresponds to a pair of different agents and each column to a pair of different items. Then, the first and the second set of constraints indicate that T is *doubly stochastic*. Here, ideally, we would like to use the famous Birkhoff-von Neumann theorem [10], which states that any doubly stochastic matrix is a convex combination of permutation matrices (i.e., square binary matrices with exactly one 1 at each row and each column) and conclude that the extreme solutions of the relaxation of the linear program (5) are integral and, hence, correspond to perfect matchings. Unfortunately, the linear program (5) has the additional symmetry constraint that does not allow for such a use of the Birkhoff-von Neumann theorem.

Fortunately, we can use an extension due to Cruse [19] which applies to centro-symmetric matrices.

Definition 5.2. An $N \times N$ matrix $T = (T_{u,v})_{u,v \in [N]}$ is called centro-symmetric if it satisfies $T_{u,v} = T_{N+1-u, N+1-v}$ for all $u, v \in [N]$.

THEOREM 5.3 (CRUSE [19]). *If N is even, then any $N \times N$ centro-symmetric doubly stochastic matrix is the convex combination of centro-symmetric permutation matrices.*

We use Theorem 5.3 in the proof of the next lemma.

LEMMA 5.4. *Any extreme solution of the relaxation of the linear program (5) is integral.*

PROOF. We will define an alternative representation of a feasible solution \mathbf{t} of the relaxation of the linear program (5) as a doubly stochastic matrix T with $N = n(n-1)$ rows and columns. To do so, we will use a particular mapping of each pair of different agents (respectively, of each pair of different items) to particular rows (respectively, columns) of the matrix T . This particular mapping will allow us to argue that the matrix T is centro-symmetric. As N is even, Theorem 5.3 will give us that T is a convex combination of centro-symmetric permutation matrices, which correspond to integral solutions.

As both sets \mathcal{N} and \mathcal{I} contain n elements each, we may view them as integers from $[n]$. We define the bijection π from ordered pairs of different integers from $[n]$ to integers of $[N]$ as follows. For every ordered pair (i, k) of different integers from $[n]$, let

$$\pi(i, k) = \sum_{h=1}^{i-1} (n-h) + k - i$$

if $i < k$, and

$$\pi(i, k) = n(n-1) + 1 - \pi(k, i)$$

otherwise. By this definition, we have

$$\pi(i, k) + \pi(k, i) = n(n-1) + 1. \quad (7)$$

Note that, for $i = 1, 2, \dots, n-1$ and $k = i+1, \dots, n$, $\pi(i, k)$ takes all distinct integer values from 1 to $n(n-1)/2$. Then, for the remaining pairs (i, k) with $i = 2, \dots, n$ and $k = 1, \dots, i-1$, $\pi(i, k)$ takes all distinct integer values from $n(n-1)$ down to $n(n-1)/2 + 1$. Hence, since $N = n(n-1)$, each

distinct ordered pair of different integers from $[n]$ is mapped to a different integer of $[N]$ under π . Hence, π is indeed a bijection. Now, for every quadruple $(i, j, k, \ell) \in \mathcal{X}$, we store the value of $t(i, j, k, \ell)$ in the entry $T_{\pi(i,k), \pi(j,\ell)}$ of matrix T . By the properties of π , the matrix T is well-defined.

We will complete the proof by showing that T is centro-symmetric. Indeed, let u and v be any integers in $[N]$ and assume that $u = \pi(i, k)$ and $v = \pi(j, \ell)$ for pairs of distinct integers (i, k) and (j, ℓ) . We have

$$\begin{aligned} T_{u,v} &= T_{\pi(i,k), \pi(j,\ell)} = t(i, j, k, \ell) = t(k, \ell, i, j) \\ &= T_{\pi(k,i), \pi(\ell,j)} = T_{N+1-\pi(i,k), N+1-\pi(j,\ell)} = T_{N+1-u, N+1-v}, \end{aligned}$$

i.e., T is indeed centro-symmetric. The first and sixth equalities follow by the definition of u and v . The second and fourth equalities follow by the definition of the entries of matrix T . The third equality is the symmetry constraint of linear program (5). The fifth equality follows by (7). \square

Hence, the execution of the ellipsoid algorithm on the relaxation of the linear program (5) will return an integral solution that corresponds to a solution of 2EBM. The next statement summarizes the above discussion.

THEOREM 5.5. *2EBM can be solved in polynomial time.*

We are ready to show how solutions to appropriately defined instances of 2EBM can be used as separation oracles for solving the linear program (3) when SW is the utilitarian (Section 5.1), egalitarian (Section 5.2), and log-Nash (Section 5.3) social welfare.

5.1 Utilitarian Social Welfare

By the definition of the utilitarian social welfare, we have

$$U(b) = \sum_{(i,j) \in b} v_i(j) = \sum_{(i,j) \in b} \sum_{\substack{(k,\ell) \in b: \\ k \neq i}} \frac{v_i(j) + v_k(\ell)}{2(n-1)},$$

and, using $SW(b) = U(b)$, the constraint of the dual linear program (3) corresponding to a matching $b \in \mathcal{M}$ is equivalent to

$$\sum_{(i,j) \in b} \sum_{\substack{(k,\ell) \in b: \\ k \neq i}} \left((v_i(j) - v_i(\ell)) \cdot y(i, j, k) + \frac{v_i(j) + v_k(\ell)}{2(n-1)} + \frac{z}{n(n-1)} \right) \leq 0. \quad (8)$$

So, consider the instance of 2EBM with edge weights defined as

$$\psi(i, j, k, \ell) = (v_i(j) - v_i(\ell)) \cdot y(i, j, k) + \frac{v_i(j) + v_k(\ell)}{2(n-1)} + \frac{z}{n(n-1)}.$$

Then, for a matching $b \in \mathcal{M}$, we have that the objective function of 2EBM, $\Psi(b)$, shown in (4), is equal to the left-hand-side of inequality (8) and, consequently, to the left-hand-side of the constraint of the dual linear program (3), when the utilitarian definition of the social welfare is used.

Now, the separation oracle for the dual linear program (3) works as follows. It solves the instance of 2EBM just described and computes a matching $b^* \in \mathcal{M}$ that maximizes the quantity $\Psi(b)$, i.e., $b^* \in \arg \max_{b \in \mathcal{M}} \Psi(b)$. If $\Psi(b^*) > 0$, the constraint corresponding to the matching b^* in the dual linear program (3) is returned as a violating constraint. Otherwise, it must be $\Psi(b) \leq 0$ for every matching $b \in \mathcal{M}$ and the separation oracle correctly returns that no such violating constraint exists.

5.2 Egalitarian Social Welfare

Let L denote the different values the valuations $v_i(j)$ of an agent i for item j can get, i.e., $L = \{v_i(j) : i \in \mathcal{N}, j \in \mathcal{I}\}$. For $e \in L$, denote by \mathcal{M}_e the set of perfect matchings so that for any agent i that is assigned to item j , it holds that $v_i(j) \geq e$. Observe that the perfect matching $b \in \mathcal{M}$ belongs to set \mathcal{M}_e for every $e \leq E(b)$. Then, the constraints of the dual linear program (3) for the egalitarian definition of the social welfare are captured by the following set of constraints:

$$\sum_{(i,j) \in b} \sum_{\substack{(k,\ell) \in b: \\ k \neq i}} \left((v_i(j) - v_i(\ell)) \cdot y(i, j, k) + \frac{e + z}{n(n-1)} \right) \leq 0, \quad e \in L, b \in \mathcal{M}_e \quad (9)$$

Indeed, for every matching $b \in \mathcal{M}$, the set of constraints (9) contains the constraint corresponding to b in the dual linear program (3) with $\text{SW} = E$ and, possibly, the redundant constraints

$$z + \sum_{\substack{(i,j) \in b \\ k \in \mathcal{N} \setminus \{i\}}} (v_i(j) - v_i(b(k))) \cdot y(i, j, k) + e \leq 0,$$

for $e \in L$ with $e < E(b)$ (if any). So, to design the separation oracle for the dual linear program (3), it suffices to design a separation oracle for the set of constraints (9), for each of the $O(n^2)$ different values of $e \in L$. We now show how to do so.

For $e \in L$, let \mathcal{X}_e be the subset of \mathcal{X} such that $v_i(j) \geq e$ and $v_k(\ell) \geq e$. Essentially, the quadruples of \mathcal{X}_e correspond to all possible (ordered) pairs of different edges in a perfect matching of \mathcal{M}_e . Now, for every $e \in L$, consider the instance of 2EBM with weights

$$\psi(i, j, k, \ell) = (v_i(j) - v_i(\ell)) \cdot y(i, j, k) + \frac{e + z}{n(n-1)}$$

for quadruple $(i, j, k, \ell) \in \mathcal{X}_e$. Then, for a matching $b \in \mathcal{M}_e$, the objective function of 2EBM, $\Psi(b)$, shown in (4), is equal to the left-hand-side of inequality (9). Now, for each $e \in L$, the separation oracle computes the matching b_e^* that maximizes the quantity $\Psi(b)$ among all matchings of \mathcal{M}_e . A violating constraint (corresponding to matching b_e^*) is then found if $\Psi(b_e^*) > 0$ for some $e \in L$. Otherwise, the separation oracle concludes that no constraint is violated.

5.3 Log-Nash Social Welfare

Observe that the definition of the log-Nash social welfare implies

$$\lg N(b) = \sum_{(i,j) \in b} \ln v_i(j) = \sum_{(i,j) \in b} \sum_{\substack{(k,\ell) \in b: \\ k \neq i}} \frac{\ln v_i(j) + \ln v_k(\ell)}{2(n-1)}.$$

Hence, the only modification that is required in the approach we followed for the utilitarian social welfare is to replace the term $\frac{v_i(j) + v_k(\ell)}{2(n-1)}$ by $\frac{\ln v_i(j) + \ln v_k(\ell)}{2(n-1)}$ in the definition of ψ .

6 INTERIM ENVY-FREE ALLOCATIONS WITH PAYMENTS

In this section, we extend the notion of interim envy-freeness by accompanying lotteries over allocations with payments to/from the agents. In this case, the definition of iEF uses both the value an agent has for item bundles and the payment she receives or contributes. We distinguish between three different types of payments. A vector of agent-dependent payments or *A-payments* consists of a payment p_i for each agent $i \in \mathcal{N}$. More refined payments are defined using additional information for an allocation instance. We say that the agents receive *bundle-dependent payments* or *B-payments* when each agent is associated with a payment of $p(S)$ when she receives the bundle of items S .

Finally, we say that the agents get *allocation-dependent payments* or C-payments when each agent i is associated with payment $p_i(A)$ in allocation A .

We now extend the notion of interim envy-freeness to pairs of lotteries and payments by distinguishing between the three payment types.

DEFINITION 6.1. *We say that a pair of a lottery \mathbf{Q} and a vector of A-payments \mathbf{p} is iEF if for every pair of agents $i, k \in \mathcal{N}$ and every bundle of items S agent i can get under \mathbf{Q} , it holds $v_i(S) + p_i \geq \mathbb{E}_{A \sim \mathbf{Q}}[v_i(A_k) | A_i = S] + p_k$.*

DEFINITION 6.2. *We say that a pair of a lottery \mathbf{Q} and a vector of B-payments \mathbf{p} per bundle of items is iEF if for every pair of agents $i, k \in \mathcal{N}$ and every bundle of items S agent i can get under \mathbf{Q} , it holds $v_i(S) + p(S) \geq \mathbb{E}_{A \sim \mathbf{Q}}[v_i(A_k) + p(A_k) | A_i = S]$.*

For C-payments, we give a more general definition that allows for marginal violations of iEF. The notion of ϵ -iEF will be useful later in Section 6.3.

DEFINITION 6.3. *Let $\epsilon \geq 0$. We say that a pair of a lottery \mathbf{Q} and a vector of C-payments \mathbf{p} per agent and allocation is ϵ -iEF if for every pair of agents $i, k \in \mathcal{N}$ and every bundle of items S agent i can get under \mathbf{Q} , it holds $v_i(S) + \mathbb{E}_{A \sim \mathbf{Q}}[p_i(A) | A_i = S] \geq \mathbb{E}_{A \sim \mathbf{Q}}[v_i(A_k) + p_k(A) | A_i = S] - \epsilon$.*

The term iEF with C-payments is used alternatively to 0-iEF. We remark that the payments are added to the value agents have for bundles in the above definitions. So, in general, the payments are assumed to be received by the agents. To represent payments that are contributed by the agents, we can allow negative entries in the payment vectors. We also remark that the definitions refer to general allocation instances. Indeed, the result we present in Section 6.1 applies to general instances. Then, in Sections 6.2 and 6.3, we restrict our attention to matching instances and adapt the definitions accordingly.

Our technical contribution regarding iEF allocations with payments is many-fold. First, we characterize in Section 6.1 those lotteries \mathbf{Q} that can be complemented with vectors of A-payments \mathbf{p} , so that the pair \mathbf{Q}, \mathbf{p} is iEF with A-payments. There, our focus is on the existence of payments, with no additional restrictions on them. In Sections 6.2 and 6.3, we specifically consider two particular optimization problems that involve iEF allocations with payments; we define these problems in the following.

In both optimization problems, we are given an allocation instance and the objective is to compute a lottery \mathbf{Q} over allocations and a payment vector \mathbf{p} so that the pair \mathbf{Q}, \mathbf{p} is iEF with payments. In *subsidy minimization*, the payments are subsidies given to the agents by an external authority. So, the corresponding entries in the payment vector \mathbf{p} are constrained to be non-negative. The goal of subsidy minimization is to find an iEF allocation and accompanying payments, such that the total expected amount of subsidies is minimized; the objective is equal to $\sum_{i \in \mathcal{N}} p_i$, $\mathbb{E}_{A \sim \mathbf{Q}}[\sum_{i \in \mathcal{N}} p(A_i)]$, or $\mathbb{E}_{A \sim \mathbf{Q}}[\sum_{i \in \mathcal{N}} p_i(A)]$, depending on whether \mathbf{p} is an A-, B-, or C-payment vector, respectively. Subsidy minimization is the generalization of the problem that was recently studied for deterministic allocations and envy-freeness in [12, 15, 27].

Our second optimization problem is called *utility maximization* and can be thought of as an extension of rent division [24] to lotteries and iEF. There is a rent R and the payments are contributions from the agents that, in expectation, should sum up to R . So, the entries in the payment vector \mathbf{p} are constrained to be non-positive. The goal is to find an iEF allocation and accompanying payments, such that the minimum expected utility over all agents is maximized. Depending on whether the problem asks for A-, B-, or C-payments, the utility of agent $i \in \mathcal{N}$ from allocation A is $v_i(A_i) + p_i$, $v_i(A_i) + p(A_i)$, and $v_i(A_i) + p_i(A)$, respectively.

As we will see in Section 6.2, both subsidy minimization and utility maximization admit much better solutions compared to their versions with deterministic allocations and envy-freeness that

had been previously studied in the literature. In addition, the quality of solutions depends on the type of payments. We demonstrate that there is no general advantage of A- or B-payments against each other; this justifies the importance of both types of payments. Clearly, C-payments allow for the best possible solutions as they generalize both A- and B-payments. In Section 6.3, we restrict our focus on matching instances and show how to solve subsidy minimization and utility maximization efficiently, by exploiting the machinery we developed in Section 5.

6.1 A Characterization for A-Payments

We now extend the notion of envy-freeability from recent works focusing on the use of subsidies in fair division settings (e.g., see [12, 15, 27]), and earlier studies in rent division (e.g., see [4, 35]). Given an allocation instance and a lottery \mathbf{Q} over allocations, we say that \mathbf{Q} is *iEF-able* with A-payments if there is a vector \mathbf{p} of A-payments so that the pair \mathbf{Q}, \mathbf{p} is iEF. Even though we will not need these terms here, we can define the term iEF-ability with B- or C-payments analogously.

Our main result in this section (Theorem 6.5) will be a characterization of the lotteries that are iEF-able with A-payments. The notion of the *interim envy graph* will be very useful; it extends the notion of the envy-graph that is central in the characterization of envy-freeable allocations (see, e.g., [27]).

DEFINITION 6.4. *Given a lottery \mathbf{Q} over allocations of the items in set \mathcal{I} to the agents in set \mathcal{N} , the interim envy graph $\text{iEG}(\mathbf{Q}, \mathcal{N}, \mathcal{I})$ is a complete directed graph with n nodes corresponding to the agents of \mathcal{N} , and edge weights defined as*

$$w(i, k) = \max_{\substack{S \subseteq \mathcal{I}: \\ \Pr_{A \sim \mathbf{Q}}[A_i = S] > 0}} \{ \mathbb{E}_{A \sim \mathbf{Q}}[v_i(A_k) | A_i = S] - v_i(S) \}$$

for every directed edge (i, k) .

Our characterization follows; it extends well-known characterizations for deterministic envy-freeable allocations, e.g., see [24, 27]. A quick comparison reveals that the second condition in Theorem 6.5 is much less restrictive than an analogous condition for envy-freeability, which asserts that envy-freeable allocations locally maximize the utilitarian social welfare among all possible redistributions of the bundles to the agents. This justifies our claim that the space of iEF-able lotteries is quite rich.

THEOREM 6.5. *For a lottery \mathbf{Q} over allocations, the following statements are equivalent:*

- (i) \mathbf{Q} is iEF-able with A-payments.
- (ii) For every agent $i \in \mathcal{N}$, let S_i be any bundle of items that is allocated to agent i with positive probability according to \mathbf{Q} . Also, let $\sigma : \mathcal{N} \rightarrow \mathcal{N}$ be any permutation of agents. Then,

$$\sum_{i \in \mathcal{N}} v_i(S_i) \geq \sum_{i \in \mathcal{N}} \mathbb{E}_{A \sim \mathbf{Q}}[v_i(A_{\sigma(i)}) | A_i = S_i].$$

- (iii) The interim envy graph $\text{iEG}(\mathbf{Q}, \mathcal{N}, \mathcal{I})$ has no cycle of positive weight.

6.2 Contrasting A-Payments with B-Payments

We now attempt a comparison between different types of payments. First, we remark that iEF lotteries with A-payments (or B-payments) can yield much lower total expected subsidies and much higher minimum expected utility for utility maximization instances, compared to envy-free allocations with payments. This should be clear given Lemma 3.3; we give explicit bounds on subsidy minimization and utility maximization with the next example.

Example 6.6. Consider the instance at the left of Table 2. By the characterization of Halpern and Shah [27], we know that in matching instances only the most efficient allocation of items to agents is envy-freeable. Therefore, $a-b-c$ and $a-c-b$ are the only envy-freeable allocations and this is possible with a payment of $1/3$ to agent 1 (or, to item a) and no payments to the other two agents (or, to the other two items). In contrast, the lottery that has both allocations in its support with equal probability is iEF without any payments.

	a	b	c
1	$1/3$	$2/3$	0
2	0	$1/2$	$1/2$
3	0	$1/2$	$1/2$

	a	b	c
1	$1/4$	$3/4$	0
2	0	$1/2$	$1/2$
3	0	$1/2$	$1/2$

Table 2. An instance of subsidy minimization (left) and utility maximization (right) with three agents where envy-freeness is inferior to iEF with A- and B-payments.

Now, consider the matching instance at the right of Table 2 and let $R = 1$. Again, $a-b-c$ and $a-c-b$ are the only envy-freeable allocations. The rent shares that make each of them EF are 0, $1/2$, and $1/2$ to agents 1, 2, and 3, respectively. Note that agents 2 and 3 obtain utility 0 under these payments. In contrast, the lottery that has both allocations in its support with equal probability is iEF with payments $1/6$, $5/12$, and $5/12$ by agents 1, 2, and 3, (or, to items a , b , c) which sum up to 1. The (expected) utility of each agent is then $1/12$. \square

We now compare A-payments to B-payments in terms of the quality of the solutions they yield the two optimization problems. In particular, the proof of Theorem 6.7 presents instances where B-payments are superior to A-payments.

THEOREM 6.7. *The solution of subsidy minimization and utility maximization with B-payments can be strictly better than the solution of the corresponding problems with A-payments.*

The proof of Theorem 6.7 shows two instances, of subsidy minimization and utility maximization, respectively. For subsidy minimization, the expected amount of subsidies achieved with B-payments is arbitrarily close to 0, while A-payments need a constant amount of subsidies. Similarly, in the utility maximization instance, the minimum expected agent utility is arbitrarily close to 0 with A-payments and considerably higher with B-payments.

Perhaps surprisingly, there are also instances where A-payments are preferable to B-payments.

THEOREM 6.8. *The solution of subsidy minimization and utility maximization with A-payments can be strictly better than the solution of the corresponding problems with B-payments.*

6.3 Computing ϵ -iEF Allocations with C-Payments

In this section, we show how to solve efficiently subsidy minimization and utility maximization when we are allowed to use C-payments (and sharp approximations of iEF). Our algorithms use linear programming and the machinery we developed in Section 5. Our result for subsidy minimization is the following.

THEOREM 6.9. *Let $\epsilon > 0$. Consider an instance of subsidy minimization with C-payments and let OPT be the value of its optimal solution. Our algorithm computes a lottery \mathbf{Q} and a C-payment vector \mathbf{p} of expected value OPT , so that the pair \mathbf{Q}, \mathbf{p} is ϵ -iEF with C-payments.*

We begin by defining a linear program for computing an iEF pair of lottery and C-payment vector. We use the variable vector \mathbf{x} to denote the lottery. By Definition 6.3, the iEF constraint for agent $i \in \mathcal{N}$ who is assigned item $j \in \mathcal{I}$ with positive probability under lottery \mathbf{x} and another agent $k \in \mathcal{N} \setminus \{i\}$ is

$$v_i(j) + \mathbb{E}_{b \sim \mathbf{x}}[p_i(b)|b(i) = j] \geq \mathbb{E}_{b \sim \mathbf{x}}[v_i(b(k)) + p_k(b)|b(i) = j]. \quad (10)$$

By multiplying the left-hand-side of (10) with the positive probability $\Pr_{b \sim \mathbf{x}}[b(i) = j]$, we get

$$\begin{aligned} & v_i(j) \cdot \Pr_{b \sim \mathbf{x}}[b(i) = j] + \mathbb{E}_{b \sim \mathbf{x}}[p_i(b)|b(i) = j] \cdot \Pr_{b \sim \mathbf{x}}[b(i) = j] \\ &= \sum_{b \in \mathcal{M}_{ij}} x(b) \cdot v_i(j) + \sum_{b \in \mathcal{M}_{ij}} x(b) \cdot p_i(b). \end{aligned} \quad (11)$$

By multiplying the right-hand-side of (10) again with $\Pr_{b \sim \mathbf{x}}[b(i) = j]$, we obtain

$$\begin{aligned} & \mathbb{E}_{b \sim \mathbf{x}}[v_i(b(k)) + p_k(b)|b(i) = j] \cdot \Pr_{b \sim \mathbf{x}}[b(i) = j] \\ &= \sum_{b \in \mathcal{M}_{ij}} x(b) \cdot v_i(b(k)) + \sum_{b \in \mathcal{M}_{ij}} x(b) \cdot p_k(b). \end{aligned} \quad (12)$$

Hence, using (11) and (12), (10) yields

$$\sum_{b \in \mathcal{M}_{ij}} (x(b) \cdot (v_i(j) - v_i(b(k))) + x(b) \cdot p_i(b) - x(b) \cdot p_k(b)) \geq 0. \quad (13)$$

Notice the products $x(b) \cdot p_i(b)$ and $x(b) \cdot p_k(b)$ in the above expression. In such terms, both factors are unknowns that we have to compute. As $p_i(b)$ always appears multiplied with $x(b)$ in the above expressions, we can avoid non-linearities by introducing the variable $t_i(b)$ for every agent $i \in \mathcal{N}$ and matching $b \in \mathcal{M}$ to be thought of as equal to $x(b) \cdot p_i(b)$. With this interpretation in mind, equation (13) becomes

$$\sum_{b \in \mathcal{M}_{ij}} (x(b) \cdot (v_i(j) - v_i(b(k))) + t_i(b) - t_k(b)) \geq 0.$$

Furthermore, our objective is to minimize $\sum_{b \in \mathcal{M}} \sum_{i \in \mathcal{N}} t_i(b)$ since

$$\sum_{b \in \mathcal{M}} \sum_{i \in \mathcal{N}} t_i(b) = \sum_{b \in \mathcal{M}} \sum_{i \in \mathcal{N}} x(b) \cdot p_i(b) = \mathbb{E}_{b \sim \mathbf{x}} \left[\sum_{i \in \mathcal{N}} p_i(b) \right].$$

Summarizing, our linear program for subsidy minimization is

$$\begin{aligned} & \text{minimize} && \sum_{b \in \mathcal{M}} \sum_{i \in \mathcal{N}} t_i(b) \\ & \text{subject to} && \sum_{b \in \mathcal{M}_{ij}} (x(b) \cdot (v_i(j) - v_i(b(k))) + t_i(b) - t_k(b)) \geq 0, i \in \mathcal{N}, j \in \mathcal{I}, k \in \mathcal{N} \setminus \{i\} \\ & && \sum_{b \in \mathcal{M}} x(b) = 1 \\ & && x(b) \geq 0, && b \in \mathcal{M} \\ & && t_i(b) \geq 0, && b \in \mathcal{M}, i \in \mathcal{N} \end{aligned} \quad (14)$$

LEMMA 6.10. *The linear program (14) can be solved in polynomial time.*

The proof of Lemma 6.10 follows along similar lines to the approach we followed in Section 5. A solution of the linear program (14) naturally gives an iEF pair of lottery \mathbf{x} and C-payments vector \mathbf{p} when $x(b) = 0$ for a matching $b \in \mathcal{M}$ implies that $t_i(b) = 0$ for every agent $i \in \mathcal{N}$. Unfortunately, we have not excluded the case that $x(b) = 0$ and $t_i(b) > 0$ in the solution of the linear program

(14). We take care of such cases by modifying the solution returned by our algorithm and violating the iEF condition marginally. The next lemma completes the proof of Theorem 6.9.

LEMMA 6.11. *For every $\epsilon > 0$, given any extreme solution to the linear program (14) of objective value OPT , an ϵ -iEF lottery \mathbf{x}' with C-payments \mathbf{p}' of total expected value OPT can be computed in polynomial time.*

Our result for utility maximization is the following.

THEOREM 6.12. *Let $\epsilon > 0$. Consider an instance of utility maximization with C-payments and let OPT be the value of its optimal solution. Our algorithm computes a lottery \mathbf{Q} and a C-payment vector \mathbf{p} of expected value at least $OPT - \epsilon$, so that the pair \mathbf{Q}, \mathbf{p} is ϵ -iEF with C-payments.*

7 OPEN PROBLEMS

We believe that interim envy-freeness can be a very influential fairness notion and can play for lotteries of allocations the role that envy-freeness has played for deterministic allocations. Our work aims to reinvigorate the study of this notion; we hope this will further intensify the study of fairness in random allocations overall. At the conceptual level, iEF can inspire new fairness notions, analogous to known relaxations of envy-freeness, such as envy-freeness up to one (EF1; see [13]) and up to any item (EFX; see [17]), that have become very popular recently. Even though it is very tempting, we refrain from proposing additional definitions here.

An appealing feature of iEF lotteries is that they are efficiently computable in matching allocation instances. Of course, besides the importance of the ellipsoid algorithm in theory (see, e.g., [25, 32]), our methods have apparent limitations. Combinatorial algorithms for solving the computational problems addressed in Sections 5 and 6.3 are undoubtedly desirable. An intermediate first step would be to design a combinatorial algorithm for the maximum edge-pair-weighted perfect bipartite matching problem (2EBM). This could be useful elsewhere, as 2EBM is a very natural problem with possible applications in other contexts.

At the technical level, there is room for several improvements of our results. Our bounds on the price of iEF with respect to the average Nash social welfare have a gap between $\Omega(\sqrt{n})$ (Theorem 4.5) and $O(n)$ (Theorem 4.2). Also, in Section 5, we show how to compute iEF lotteries that maximize the expected log-Nash social welfare. Although maximizing log-Nash and average Nash social welfare are equivalent goals for deterministic allocations, this is not the case for lotteries. Computing iEF lotteries of maximum expected average Nash welfare is elusive at this point.

We left for the end the many open problems that are related to iEF with payments. Our characterization in Section 6.1 has been used only in the proof of Theorem 6.7. It would be interesting to see whether it has wider applicability and, in particular, whether it can lead to efficient algorithms for computing iEF pairs of lotteries with A-payments. This is not clear to us, as our characterization seems to be much less restrictive than existing ones for envy-freeability (which, e.g., have given rise to transforming the rent division problem to a bipartite matching computation [24]). Furthermore, characterizations of iEF-ability with B- or C-payments are currently elusive. Our results in Section 6.2 reveal gaps on the quality of solutions for the two optimization problems that the different types of payments allow. Our analysis in the proofs of Theorems 6.7 and 6.8 is not tight; determining the maximum gap between A-, B-, and C-payments on the quality of solutions for subsidy minimization and utility maximization is open. Finally, from the computational point of view, our solutions to subsidy minimization and utility maximization yield pairs of lotteries and C-payments that are only approximately iEF. Can these problems be solved exactly? Also, solving both subsidy minimization and utility maximization with A- or B-payments will be of practical importance. What is the complexity of these problems?

ACKNOWLEDGMENTS

This work was partially supported by the European Commission under COST Action 16228 – European Network for Game Theory (GAMENET).

REFERENCES

- [1] Atila Abdulkadiroğlu and Tayfun Sönmez. 1998. Random Serial Dictatorship and the Core from Random Endowments in House Allocation Problems. *Econometrica* 66, 3 (1998), 689–701.
- [2] Atila Abdulkadiroğlu, Tayfun Sönmez, and M. Utku Ünver. 2004. Room assignment-rent division: A market approach. *Social Choice and Welfare* 22, 3 (2004), 515–538.
- [3] Martin Aleksandrov, Haris Aziz, Serge Gaspers, and Toby Walsh. 2015. Online Fair Division: Analysing a Food Bank Problem. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*. 2540–2546.
- [4] Ahmet Alkan, Gabrielle Demange, and David Gale. 1991. Fair Allocation of Indivisible Goods and Criteria of Justice. *Econometrica* 59, 4 (1991), 1023–1039.
- [5] Enriqueta Aragones. 1995. A derivation of the money Rawlsian solution. *Social Choice and Welfare* 12, 3 (1995), 267–276.
- [6] Haris Aziz. 2019. A Probabilistic Approach to Voting, Allocation, Matching, and Coalition Formation. *The Future of Economic Design* Springer, 45–50 (2019).
- [7] Haris Aziz. 2020. Simultaneously Achieving Ex-ante and Ex-post Fairness. In *Proceedings of the 16th International Conference on Web and Internet Economics (WINE) (Lecture Notes in Computer Science)*, Vol. 12495. 341–355.
- [8] Haris Aziz and Florian Brandl. 2020. The Vigilant Eating Rule: A General Approach for Probabilistic Economic Design with Constraints. *CoRR* abs/2008.08991 (2020).
- [9] Dimitris Bertsimas, Vivek F. Farias, and Nikolaos Trichakis. 2011. The Price of Fairness. *Operations Research* 59, 1 (2011), 17–31.
- [10] Garrett Birkhoff. 1946. Three observations on linear algebra (in spanish). *Univ. Nac. Tacuman, Rev. Ser. A* 5 (1946), 147–151.
- [11] Anna Bogomolnaia and Herve Moulin. 2001. A New Solution to the Random Assignment Problem. *Journal of Economic Theory* 100, 2 (2001), 295 – 328.
- [12] Johannes Brustle, Jack Dippel, Vishnu V. Narayan, Mashbat Suzuki, and Adrian Vetta. 2020. One Dollar Each Eliminates Envy. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*, Péter Biró, Jason Hartline, Michael Ostrovsky, and Ariel D. Procaccia (Eds.). 23–39.
- [13] Eric Budish. 2011. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy* 119, 6 (2011), 1061–1103.
- [14] Eric Budish, Yeon-Koo Che, Fuhito Kojima, and Paul Milgrom. 2013. Designing Random Allocation Mechanisms: Theory and Applications. *American Economic Review* 103, 2 (2013), 585–623.
- [15] Ioannis Caragiannis and Stavros Ioannidis. 2020. Computing envy-freeable allocations with limited subsidies. *CoRR* abs/2002.02789 (2020).
- [16] Ioannis Caragiannis, Christos Kaklamani, Panagiotis Kanellopoulos, and Maria Kyropoulou. 2012. The Efficiency of Fair Division. *Theory of Computing Systems* 50, 4 (2012), 589–610.
- [17] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. 2019. The Unreasonable Fairness of Maximum Nash Welfare. *ACM Transactions on Economics and Computation* 7, 3 (2019), 12:1–12:32.
- [18] Yann Chevaleyre, Ulle Endriss, and Nicolas Maudet. 2017. Distributed fair allocation of indivisible goods. *Artificial Intelligence* 242 (2017), 1–22.
- [19] Allan B. Cruse. 1977. Some Combinatorial properties of centrosymmetric matrices. *Linear Algebra Appl.* 16, 1 (1977), 65–77.
- [20] Geoffroy de Clippel. 2008. Equity, envy and efficiency under asymmetric information. *Economics Letters* 99 (2008), 65–67.
- [21] Duncan K. Foley. 1967. Resource allocation and the public sector. *Yale Economic Essays* 7, 1 (1967), 45–98.
- [22] Rupert Freeman, Nisarg Shah, and Rohit Vaish. 2020. Best of Both Worlds: Ex-Ante and Ex-Post Fairness in Resource Allocation. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*. 21–22.
- [23] Thibault Gajdos and Jean-Marc Tallon. 2002. Fairness under Uncertainty. *Economics Bulletin* 4 (11 2002), 1–7.
- [24] Ya'akov (Kobi) Gal, Moshe Mash, Ariel D. Procaccia, and Yair Zick. 2017. Which Is the Fairest (Rent Division) of Them All? *Journal of the ACM* 64, 6 (2017), 39:1–39:22.
- [25] Martin Grötschel, László Lovász, and Alexander Schrijver. 1988. *The Ellipsoid Method*. Springer Berlin Heidelberg, 64–101.

- [26] Claus-Jochen Haake, Matthias G. Raith, and Francis Edward Su. 2002. Bidding for envy-freeness: A procedural approach to n-player fair-division problems. *Social Choice and Welfare* 19, 4 (2002), 723–749.
- [27] Daniel Halpern and Nisarg Shah. 2019. Fair Division with Subsidy. In *Proceedings of the 12th International Symposium on Algorithmic Game Theory (SAGT) (Lecture Notes in Computer Science)*, Vol. 11801. 374–389.
- [28] Aanund Hylland and Richard Zeckhauser. 1979. The Efficient Allocation of Individuals to Positions. *Journal of Political Economy* 87, 2 (1979), 293–314.
- [29] Flip Klijn. 2000. An algorithm for envy-free allocations in an economy with indivisible objects and money. *Social Choice and Welfare* 17, 2 (2000), 201–215.
- [30] Eric S. Maskin. 1987. *On the Fair Allocation of Indivisible Goods*. Palgrave Macmillan UK, 341–349.
- [31] Thomas R. Palfrey and Sanjay Srivastava. 1987. On Bayesian Implementable Allocations. *The Review of Economic Studies* 54, 2 (1987), 193–208.
- [32] Alexander Schrijver. 1986. *Theory of Linear and Integer Programming*. John Wiley & Sons, Inc.
- [33] Francis Edward Su. 1999. Rental Harmony: Sperner’s Lemma in Fair Division. *The American Mathematical Monthly* 106, 10 (1999), 930–942.
- [34] Lars-Gunnar Svensson. 1983. Large Indivisibles: An Analysis with Respect to Price Equilibrium and Fairness. *Econometrica* 51, 4 (1983), 939–954.
- [35] Lars-Gunnar Svensson. 2009. Coalitional Strategy-Proofness and Fairness. *Economic Theory* 40, 2 (2009), 227–245.
- [36] Hal R. Varian. 1974. Equity, envy, and efficiency. *Journal of Economic Theory* 9, 1 (1974), 63–91.
- [37] Róbert F. Veszteg. 2004. *Auctions, Mechanisms, and Uncertainty*. Ph.D. Dissertation. Universitat Autònoma de Barcelona.