

Knowledge, Fairness, and Social Constraints

Haris Aziz

Data61, CSIRO and UNSW
Sydney, Australia
haris.aziz@data61.csiro.com.au

Sylvain Bouveret

Univ. Grenoble-Alpes, France
sylvain.bouveret@imag.fr

Ioannis Caragiannis

University of Patras, Greece
caragian@ceid.upatras.gr

Ira Giagkousi

University of Patras, Greece
giagkousi@ceid.upatras.gr

Jérôme Lang

CNRS, U. Paris-Dauphine, PSL, France
lang@lamsade.dauphine.fr

Abstract

In the context of fair allocation of indivisible items, fairness concepts often compare the satisfaction of an agent to the satisfaction she would have from items that are not allocated to her: in particular, envy-freeness requires that no agent prefers the share of someone else to her own share. We argue that these notions could also be defined relative to the knowledge that an agent has on how the items that she does not receive are distributed among other agents. We define a family of epistemic notions of envy-freeness, parameterized by a social graph, where an agent observes the share of her neighbours but not of her non-neighbours. We also define an intermediate notion between envy-freeness and proportionality, also parameterized by a social graph. These weaker notions of envy-freeness are useful when seeking a fair allocation, since envy-freeness is often too strong. We position these notions with respect to known ones, thus revealing new rich hierarchies of fairness concepts. Finally, we present a very general framework that covers all the existing and many new fairness concepts.

1 Introduction

Envy-freeness is one of the most important fairness requirements in the classical resource allocation problem of distributing indivisible items among agents who have cardinal valuations over these items (Bouveret, Chevaleyre, and Maudet 2016; Chevaleyre et al. 2006). However, it is quite strong and not achievable in many resource allocation instances. Hence, several authors have studied a series of weaker notions, such as proportionality (which dates back to Steinhaus (1948)), maxmin (Budish 2011) and minmax fair share (Bouveret and Lemaître 2016). Which fairness notion is the most appropriate for resource allocation? A compelling objective would be to find an allocation for the strongest fairness concept possible. An “optimal” solution to this problem would then require to have relaxations of envy-freeness of varying strength available. In this paper we define and study several such families of relaxations. We introduce the first one by an example.

Ann, Bob and Chloe are PC members in the *International Conference on Everything* and are about to “bid” for papers to review. Four papers p_1, \dots, p_4 have been submitted. Ann

and Bob both value each of these papers as 10, 6, 6 and 1, respectively. Chloe values them 1, 6, 6, and 10. All three have additive preferences. Each paper must be reviewed by exactly one reviewer. Ideally, one would like to assign papers to reviewers so that none of them would prefer the share of someone else to her own share; this is known as an *envy-free* allocation. It is easy to see that no such allocation exists here: since Ann and Bob have identical preferences, they are non-*envious* of each other only if they draw the same value from their assignments: the only ways to achieve this is to assign p_2 to one of them and p_3 to the other, or no paper to any of them; in any case, they will both envy Chloe.

However, this line of reasoning implicitly assumes that *each agent knows how the items are allocated*. This is not necessarily the case in such a context. If each of the three agents ignores to whom the papers that they do not get are assigned, we can assign the papers in such a way that *everyone considers it possible that they do not envy anyone else*: assign p_1 to Ann, p_2 and p_3 to Bob and p_4 to Chloe. Ann, who does not know how p_2 , p_3 , and p_4 are assigned, considers it possible that one of Bob and Chloe gets p_2 and the other one gets both p_3 and p_4 , in which case she does not envy anyone; such a line of reasoning also applies to Bob and to Chloe. Although this allocation is not *envy-free*, it is *epistemically envy-free*.¹

While classical *envy-freeness* is based on the assumption that the knowledge of the allocation among the agents is *maximal* (everyone knows which items every other agent gets), *epistemic envy-freeness* implicitly assumes that this knowledge is *minimal* and each agent knows only which of the items have been allocated to her. We introduce intermediate notions where an agent knows the bundles of items that a specific subset of agents get. For instance, an agent may have knowledge only of the bundles of items allocated to her and her acquaintances (or *neighbors*) in a social network. Still, she may reason about the possible allocations of the remaining items to the agents she is not connected to. We will consider such notions of *epistemic fairness* which involve *social graphs* (to model social acquaintances); clearly, the denser the graph, the stronger the notion.

¹This does not mean that every agent *considers it possible that the allocation is envy-free*. For this, each agent should have additional knowledge of the others’ preference.

As a fairness concept, envy-freeness is *knowledge-sensitive*. It explicitly compares an agent’s bundle of items to the bundles of other agents. In contrast, proportionality, which requires that the value of each agent for the items in her bundle is at least as high as her total value for all items divided by the number of agents, is not knowledge-sensitive in this sense. An agent can tell whether her bundle is proportional or not without knowing how the items she did not get are allocated. Still, social graphs can be used to define a family of notions that are intermediate between envy-freeness and proportionality. Our related fairness notion requires that each agent is non-jealous of all her neighbours and gets a proportional share compared to her non-neighbours.

As another conceptual contribution, we propose an *abstract framework* that can be used to define novel fairness notions from very basic ones. All notions of epistemic fairness that we present (and many more) can be obtained as part of this framework. On the technical level, our main objective is to relate the new fairness notions to well-known ones in terms of strength. Our results indicate that epistemic envy-freeness is weaker than envy-freeness but stronger than min-max fair share. Our extensions of epistemic envy-freeness using social graphs define a very rich hierarchy of fairness concepts that “occupy the space” between epistemic envy-freeness and envy-freeness. A similar hierarchy is obtained when combining envy-freeness, proportionality, and social graphs; these fairness concepts lie between proportionality and envy-freeness.

The rest of the paper is structured as follows. We conclude this section with a presentation and discussion of related work. The exposition of our technical results begins with notation and background definitions in Section 2. In Section 3, we define epistemic envy-freeness and relate it to other known fairness notions. The definition of epistemic envy-freeness concepts that involve social graphs and the results establishing their hierarchy are presented in Section 4. Section 5 is devoted to our novel fairness concept that combines proportionality, envy-freeness, and social graphs. In Section 6 we present our abstract framework and demonstrate how the several well-known and new fairness concepts can be obtained. We conclude in Section 7 with directions for future research.

Related work. As envy-freeness is not always achievable, several relaxations of it have been studied in the literature. These include the fairness concepts of envy-freeness up to some item (EF1) and up to any single item (EFX) (Caragiannis et al. 2016; Lipton et al. 2004; Plaut and Roughgarden 2017, for instance). The distance from envy-freeness (quantified as the envy ratio or degree of envy) has also been used as an objective to minimize (Lipton et al. 2004; Nguyen and Rothe 2014).

Chen and Shah (2017) revisit envy-freeness by assuming that allocations are not public knowledge (to the agents). At first glance, this seems to be conceptually related to our notions of epistemic fairness. Both the goal and the technical content in the two papers are very different, though. Instead of defining a relaxation of envy-freeness, the main goal of Chen and Shah (2017) is to study the implications of hiding information from the agents (this might be desirable

for other reasons; e.g., for privacy) on the way they quantify envy. In their setting, each agent has a prior distribution about how the other agents value the items, which she can use to infer information about their bundles and to compute her expected value for them. The definition of our epistemic envy-freeness (and its variants) does not use any such prior.

Several publications make connection of fair allocations to some kind of graph or network. For instance, some papers (Abebe, Kleinberg, and Parkes 2017; Bei, Qiao, and Zhang 2017; Chevaleyre, Endriss, and Maudet 2017) assume an underlying social network over the agents and define envy-freeness or proportionality constraints for each agent in relation to her neighbors only. Unlike in our fairness concepts, the items that are allocated to non-neighbors are completely ignored in this way.

In very recent work, the resource allocation instance is augmented by a graph representing structure among the items. The goal is to identify allocations in which the set of resources allocated to each agent is connected (Bouveret et al. 2017). This is a natural extension of the requirement for contiguous allocations in the cake-cutting model. See the books (Brams and Taylor 1996; Robertson and Webb 1998) for a more detailed coverage of cake-cutting and the paper of Procaccia (2016) for a recent short survey.

2 Definitions and Notation

We consider resource allocation instances that consist of n agents and a set M of items. We identify the agents by positive integers and denote their set by $[n] = \{1, 2, \dots, n\}$.

An *allocation* of the items in M to the agents of $[n]$ is a partition of M into n sets, e.g., $A = (A_1, A_2, \dots, A_n)$ with $M = \bigcup_{i \in [n]} A_i$ and $A_i \cap A_j = \emptyset$ for $i \neq j$. We will refer to the set A_i as the set (or *bundle*) of items allocated to agent i .

Each agent i has a non-negative *valuation function* $v_i : M \rightarrow \mathbb{R}_{\geq 0}$ defined over the items. Then, with some abuse of notation, the *value* of agent i for a set of items $S \subseteq M$ is $v_i(S) = \sum_{g \in S} v_i(g)$. Let us now recall some well-known notions of fairness for allocations. We begin with envy-freeness and proportionality.

Definition 1. Consider a resource allocation instance with n agents and a set of items M . An allocation $A = (A_1, \dots, A_n)$ satisfies

- proportionality if $v_i(A_i) \geq \frac{1}{n} v_i(M)$ for every $i \in [n]$, and
- envy-freeness if $v_i(A_i) \geq v_i(A_j)$ for every $i, j \in [n]$.

Alternatively, we will say that the allocation A is proportional (PROP) and envy-free (EF), respectively. We will use the same abbreviations PROP and EF to refer to proportionality and envy-freeness. We will also refer to the quantity $\frac{1}{n} v_i(M)$ as the proportionality *threshold* of agent i .

These two fairness notions have received enormous attention in the fair division literature. It is not hard to see (e.g., consider two agents and a single item) that there are resource allocation instances which do not have any PROP or EF allocation. It is also well-known that EF *implies* PROP in the sense that an EF allocation is also PROP. This implication is *strict* in the sense that, not only are there PROP allocations

that are not EF but that, more importantly, there are resource allocation instances which have some PROP allocation but no EF allocation.

Bouveret and Lemaître (2016) considered more fairness concepts and expanded the list of implications. We will use two of them here.

Definition 2. Consider a resource allocation instance with n agents and a set of items M and let \mathcal{F} denote the set of all allocations.

- The min-max fair share of agent i is defined as

$$u_i^{\text{mFS}}(M, n) = \min_{A^i \in \mathcal{F}} \max_{j \in [n]} v_i(A_j^i).$$

An allocation $A = (A_1, \dots, A_n)$ satisfies the min-max fair share criterion (or, simply, is mFS) if $v_i(A_i) \geq u_i^{\text{mFS}}(M, n)$ for every $i \in [n]$.

- The max-min fair share of agent i is defined as

$$u_i^{\text{MFS}}(M, n) = \max_{A^i \in \mathcal{F}} \min_{j \in [n]} v_i(A_j^i).$$

An allocation $A = (A_1, \dots, A_n)$ satisfies the max-min fair share criterion (or, simply, is MFS) if $v_i(A_i) \geq u_i^{\text{MFS}}(M, n)$ for every $i \in [n]$.

Again, the abbreviations mFS and MFS will be used to refer to the corresponding fairness concepts as well. Like PROP, both mFS and MFS define a threshold for each agent that depends on her valuation function and the number of agents and require that her value for the bundle of items the agent receives is at least as high as her threshold.

EF is the strongest among the four fairness concepts above. It implies mFS, which implies PROP, which in turn implies MFS. These implications are rather easy to prove; what is considerable trickier to show is that these implications are strict. For example, mFS does not imply EF as there are resource allocation instances with an mFS allocation but with no EF allocation. See the paper of Bouveret and Lemaître (2016) for nice counterexamples showing that the above implications are all strict. Interestingly, there are resource allocation instances that are not MFS (Procaccia and Wang 2014).

In this way, a five-level hierarchy of fairness concepts is defined, having EF instances on the top² level, mFS instances on the second level, and so on, while all instances are in the bottom level. In the following sections, we define several new fairness notions which, as we show, occupy intermediate levels in this hierarchy.

3 Epistemic Envy-Freeness

In the definition of PROP, mFS, and MFS, the constraints for agent i (i.e., the inequalities involving the value $v_i(A_i)$ of agent i for her bundle) in an allocation A depend only on the number n of agents, the bundle A_i that is allocated

²Even though EF is the strongest fairness notion we consider here, there are even stronger ones in the literature, such as the well-known competitive equilibrium from equal incomes (CEEI) from microeconomics. CEEI implies EF and this implication is strict (Bouveret and Lemaître 2016).

to agent i , and her valuation function over the items in M . In contrast, the EF constraints for agent i depend on the whole allocation vector, including the bundles of items that the other agents get in A . The fairness notion of epistemic envy-freeness (EEF) aims to relax EF by mitigating this dependence.

Definition 3. Consider a resource allocation instance with n agents and a set of items M . An allocation $A = (A_1, \dots, A_n)$ satisfies epistemic envy-freeness (or, simply, is EEF) if for every $i \in [n]$, there exists an allocation $A^i = (A_1^i, \dots, A_n^i)$ such that $A_i = A_i^i$ and $v_i(A_i) \geq v_i(A_j^i)$ for all $j \in [n]$.

In other words, an allocation is EEF if, for every agent i , the items that are not allocated to her can be re-allocated to the other agents so that agent i does not envy anyone after the re-allocation. Let us warm up by proving simple implications that involve EEF.

Theorem 4. EF implies EEF which implies mFS.

Proof. By definition, an EF allocation A is also EEF (using $A^i = A$ for every $i \in [n]$). We now show that an EEF allocation is also mFS, thus proving the second implication. Consider an EEF allocation A in an n -agent resource allocation instance with a set of items M . We will show that the value of every agent i for her bundle is at least her min-max fair share, i.e., $v_i(A_i) \geq u_i^{\text{mFS}}(M, n)$. Since A is EEF, there exists an allocation A^i such that $v_i(A_i) \geq \max_{j \in [n]} v_i(A_j^i)$. The proof follows by observing (see Definition 2) that the RHS of this inequality is lower-bounded by $u_i^{\text{mFS}}(M, n)$. \square

The relation of EEF to known fairness concepts is summarized in Figure 1.

$$\text{EF} \implies \text{EEF} \implies \text{mFS} \implies \text{PROP} \implies \text{MFS}$$

Figure 1: Relations of fairness concepts.

Next, in Theorems 5 and 6, we show that both implications are strict. We do so using instances with three agents since EF, EEF, mFS, and PROP are all equivalent for $n = 2$. The instance in the proof of the next statement was also used by Bouveret and Lemaître (2016) to prove that mFS does not always imply EF.

Theorem 5. There exist resource allocation instances with an EEF allocation but with no EF allocation.

Proof. Consider the instance with four items A, B, C, D and three agents with the valuations that are shown in the next table:

agent	A	B	C	D
1	10	6	6	1
2	10	6	6	1
3	1	6	6	10

The allocation in which agent 1 gets item A , agent 2 gets items B and C , and agent 3 gets item D is EEF. Indeed, from

the point of view of agent 1 who gets item A , there is the re-allocation in which agent 2 gets item B and agent 3 gets items C and D so that agent 1 does not envy the two other bundles. Agent 2 is not envious of the other agents in the current allocation anyway. Agent 3, who gets item D , is not envious after re-allocating items A and B to agent 1 and item C to agent 2. On the other hand, as observed by Bouveret and Lemaître (2016), there is no EF allocation. \square

The EEF allocation in the above proof is mFS. So, in order to show that EEF is strictly stronger than mFS (i.e., that mFS does not always imply EEF), we will use a more complicated instance³ and, correspondingly, more involved arguments.

Theorem 6. *There exist resource allocation instances with an mFS allocation but with no EEF allocation.*

Proof. Consider the instance with six items A, B, C, D, E, F and three agents with the valuations that are shown in the next table:

agent	A	B	C	D	E	F
1	9	9	9	16	15	15
2	8	8	8	20	13	13
3	8	8	8	24	12	12

The min-max fair shares are 25 for agent 1, 26 for agent 2, and 24 for agent 3. To see why, observe that the total value of agents 1 and 3 for all items is 73 and 72, respectively. Then, the minimum possible value of 25 for the min-max fair share of agent 1 is obtained by the allocation $(\{A, D\}, \{B, E\}, \{C, F\})$. Similarly, the minimum possible value of 24 for the min-max fair share of agent 3 is obtained by the allocation $(\{A, B, C\}, \{D\}, \{E, F\})$. Agent 2 has maximum value 26 for the bundles of the allocation $(\{A, B, C\}, \{D\}, \{E, F\})$; it is easy to see that this is the “most balanced” allocation that defines her min-max fair share.

Now, an mFS allocation is as follows: agent 1 gets items A, B, C which she values for 27, agent 2 gets items E, F of value 26, and agent 3 gets item D of value 24.

We now show that there is no EEF allocation. First observe that if agent 1 (similarly, for agent 2) does not get any of the items D, E , and F , then, in any re-allocation, she will be envious of the agent who gets at least two items among them. So, in an EEF allocation (if any), each of the agents 1 and 2 should receive at least one item among D, E , and F . From now on, our argument will be that there is no such allocation in which each agent gets her proportionality threshold (i.e., 25, 24, and 24, respectively). Since EEF implies PROP, in this way we will have shown that no EEF allocation exists either, completing the proof of Theorem 6.

³We remark that, even though the three agents have the same total value of 23 for all items in the instance used in the proof of Theorem 5, this is not a requirement of our model. Some of the constructions in the following have this feature (i.e., the ones in the proofs of Theorems 9 and 12) while others (i.e., in the proofs of Theorems 6, 14, and 15) do not.

First assume that all items D, E , and F are allocated to agents 1 and 2. Then, the items A, B , and C should be all allocated to agent 3 (in order to get her proportionality threshold) which will leave some of agents 1 and 2 with a value that is below her proportionality threshold.

It remains to consider the case in which each of agents 1 and 2 get exactly one among items D, E , and F . In order to get the proportionality threshold, an agent (among agents 1 and 2) needs one extra item from A, B, C if item D is allocated to her and two extra items from A, B , and C otherwise. We conclude that D should be allocated to some of the agents 1 and 2 and agent 3 is left with only one item among E and F ; this is not enough and gives her a value (of only 12) that is below her proportionality threshold. \square

4 A New Fairness Concept with Social Constraints

EF and EEF stand in the two extremes of the possible awareness levels of the agents. EEF assumes that each agent is ignorant of how the items that she does not receive are allocated to the other agents, while EF assumes that all agents have full knowledge of all bundles. In this section, we define a new fairness notion by modeling intermediate awareness levels for each agent. Besides her bundle, an agent is allowed to have full knowledge for the bundles of some particular agents only.

We model such social constraints with a directed graph called the *social graph*. Each node of a social graph corresponds to an agent and a directed edge (u, v) indicates that agent u is always aware of the items allocated to agent v . Given a social graph G , we use $\text{Neig}_G(u)$ and $\text{deg}_G(u)$ to denote the set of out-neighbors (or, simply, *neighbors*) and the out-degree of node u in graph G , respectively.

Definition 7. *Let G be an n -node social graph and consider an n -agent resource allocation instance over a set of items M . An allocation $A = (A_1, \dots, A_n)$ is G -EEF if there exists an allocation $A^i = (A_1^i, \dots, A_n^i)$ for every agent $i \in [n]$ such that $A_j = A_j^i$ for every $j \in \{i\} \cup \text{Neig}_G(i)$ and $v_i(A_i) \geq v_i(A_j^i)$ for every $j \in [n]$.*

Equivalently, the definition requires that $v_i(A_i) \geq v_i(A_j)$ for every $j \in \text{Neig}_G(i)$ (i.e., agent i is not envious for her neighbors in the original allocation) and $v_i(A_i) \geq v_i(A_j^i)$ for every $j \notin \text{Neig}_G(i)$ (i.e., agent i is not envious of her non-neighbors when the items are re-allocated).

Clearly, EEF is identical to G -EEF when G contains no edges. Also, observe that the requirement imposed by G -EEF is the same when a node has out-degree $n - 2$ or $n - 1$; in both cases, G -EEF requires the corresponding agent to be non-envious of the other agents. Hence, envy-freeness is identical to G -EEF when every node of G has out-degree at least $n - 2$.

We begin with a simple observation.

Theorem 8. *Let G and H be social graphs over the same set of nodes. If G is a subgraph of H , then H -EEF implies G -EEF.*

Proof. Consider an n agent resource allocation instance with the set of items M . Let A be an H -EEF allocation

and A^i be the re-allocation of the items with $A_j = A_j^i$ for $j \in \{i\} \cup \text{Neig}_H(i)$ that satisfies $v_i(A_i) \geq v_i(A_j^i)$ for every $i, j \in [n]$. Since G is a subgraph of H , the allocation A^i satisfies $A_j = A_j^i$ for $j \in \{i\} \cup \text{Neig}_G(i)$ as well. Hence, A is G -EEF. \square

The next theorem states that for almost every other relation between social graphs H and G , H -EEF does not necessarily imply G -EEF.

Theorem 9. *Let $G = ([n], E(G))$ and $H = ([n], E(H))$ be n -node (with $n \geq 3$) social graphs which have two nodes u and v such that $\deg_H(u) \leq n - 3$ and edge (u, v) belongs to $E(G)$ but not to $E(H)$. Then, there exists an n -agent resource allocation instance that has an H -EEF but no G -EEF allocation.*

Proof. We will prove the theorem using the following n -agent instance. Since $\deg_H(u) \leq n - 3$, let w be a node different than u and v such that the directed edge (u, w) does not belong to $E(H)$. There are n agents and $n + 1$ items. The valuation of agent u is 0, 3, 3, and 4 for the first four items and 0 for the remaining ones.⁴ The valuation of agent v is 4, 3, 3, and 0 for the first four items and 0 for the remaining ones. The valuation of agent w is 6, 2, 2, and 0 for the first four items and 0 for the remaining ones. The valuation of agent $x \in [n] \setminus \{u, v, w\}$ (if any) is 0 for the first four items and $\frac{10}{n-3}$ for each of the remaining $n - 3$ items (if any). See the next table:

agent	1	2	3	4	5	...	$n + 1$
u	0	3	3	4	.	.	.
v	4	3	3	0	.	0	.
w	6	2	2	0	.	.	.
.
x	.	0	.	.	.	$\frac{10}{n-3}$.
.

We first claim that any EEF allocation and, consequently (by Theorem 8), any H -EEF or G -EEF allocation should at least satisfy the following properties:

- agent u should get item 4,
- agent v should get items 2 and 3,
- agent w should get item 1, and
- each of the remaining $n - 3$ agents should get one of the items 5, ..., $n + 1$.

The fourth property is obvious since the remaining $n - 3$ agents have non-zero value for the $n - 3$ items 5, ..., $n + 1$. Let us now argue about the necessity of the first three properties. Observe that no allocation in which agent w does not get item 1 is EEF, since the agent would then envy the agent who gets the item in any re-allocation. Also, if agent v does not get item 1 and gets only one of the items 2 and 3, she will be envious for the agent that gets item 1 in any re-allocation. So, in any EEF allocation, agent w should get item 1 and

⁴Even though the instances in the proofs of Theorems 9, 12, 14, and 15 include zero valuations for some items, our arguments carry over when replacing 0 by a very small but strictly positive value.

agent v gets both items 2 and 3. Then, agent u should get item 4 (which is the only available item for which she has non-zero value).

We now claim that such an allocation is not G -EEF. Indeed, as G contains edge (u, v) , the requirement that agent u is not envious for the bundle allocated to agent v is not satisfied. It remains to prove that such an allocation is H -EEF. Clearly, no agent besides u is envious. For agent u , since the edges (u, v) and (u, w) do not exist in H , it suffices to show that there is a re-allocation of the items 1, 2, and 3 to agents v and w so that agent u does not envy them: giving items 1 and 2 to agent w and item 3 to agent v yields the desired re-allocation. \square

We remark that the restriction on the out-degree of node u in H in the statement of Theorem 9 is necessary. For example, consider graph H in which node u is connected to every other node besides v and G is obtained by adding edge (u, v) in H . Then, any allocation that is H -EEF is also G -EEF since, in both cases, the requirement for agent u is to be EF.

Now, observe that any pair of different n -node social graphs in which all nodes have out-degree either exactly $n - 1$ or at most $n - 3$ satisfy the condition of Theorem 9 (the requirement from an agent corresponding to a node of out-degree $n - 2$ is essentially to be non-envious to every other agent). Hence, any such graph defines a different fairness concept. A careful counting of the different sets of either $n - 1$ or at most $n - 3$ outgoing edges a node can have yields that there are $(2^{n-1} - n + 1)^n$ such social graphs. Therefore, together, Theorems 8 and 9 establish a very rich sub-hierarchy of fairness concepts between EF and EEF.

5 Relaxing Graph-EEF

We now introduce another fairness notion that includes social constraints and somehow lies between PROP and EF.

Definition 10. *Let G be an n -node social graph and consider an n -agent resource allocation instance over a set of items M . An allocation $A = (A_1, \dots, A_n)$ is G -PEF if, for every agent i , $v_i(A_i) \geq v_i(A_j)$ for every $j \in \text{Neig}_G(i)$, and $v_i(A_i) \geq \frac{1}{n-1-\deg_G(i)} \cdot \sum_{j \notin \text{Neig}_G(i)} v_i(A_j)$.*

That is, G -PEF requires each agent to be non-envious for her neighbors and above the proportionality threshold compared to her non-neighbors. So, G -PEF is in a sense less demanding than G -EEF as its definition involves no re-allocation of items. To further justify G -PEF, observe that deciding whether a given allocation is G -PEF is an easy task; in contrast, deciding whether a given allocation is G -EEF can be easily seen to be computationally hard.

Clearly, G -PEF is identical to PROP if G contains no edges and is identical to EF if each node of G has out-degree either $n - 2$ or $n - 1$. We continue by showing that, given a social graph G , G -EEF is (strictly) stronger than G -PEF. We first show that G -EEF implies G -PEF.

Theorem 11. *For every social graph G , G -EEF implies G -PEF.*

Since the agents v and w have non-zero value for a single item and the $n - 3$ agents besides u , v , and w have value for the items $4, \dots, n$ only, any H -PEF or G -PEF allocation should at least satisfy the following properties:

- agents u , v , and w get items 1, 2, and 3, respectively, and
- each of the remaining $n - 3$ agents gets one of the items $4, \dots, n$.

Such an allocation is indeed H -PEF since agent u gets a value of 1 from item 1 and her total value for the $n - 1 - \deg_H(u)$ items that are allocated to agents v , w , and the other $n - 3 - \deg_H(u)$ non-neighbors of u among the last $n - 3$ agents is exactly $n - 1 - \deg_H(u)$.

In contrast, this allocation is not G -PEF. If (u, w) belongs to $E(G)$ agent u will envy the item allocated to agent w . Otherwise, the total value agent u has for the items allocated to her, her $n - 3 - \deg_G(u)$ non-neighbors among the last $n - 3$ agents, and agent w is exactly $n - \deg_G(u)$ and, hence, her value for item 1 which is allocated to her is strictly below the proportionality threshold. \square

We conclude this section by stating that G -PEF is much stronger than the notion of graph envy-freeness⁵ from Chevaleyre, Endriss and Maudet (2017). Recall that an allocation is G -EF if no agent is envious for her neighbors; clearly, G -PEF implies G -EF.

Theorem 15. *Let $G = ([n], E(G))$ be a non-complete n -node (with $n \geq 2$) social graph. Then, there exists an n -agent resource allocation instance that has a G -EF allocation but no G -PEF allocation.*

Proof. First observe that, for every graph, G -PEF implies PROP (by Theorem 13). We will prove the theorem by showing that G -EF and PROP may not be compatible. Since G is not complete, let v be a node such that $(u, v) \notin E(G)$. Then, consider the following instance. There are n items. The valuation of agent u is 1 and n for the first two items and 0 for the remaining ones. The valuation of agent v is 1 for item 2 and 0 for the remaining items. Every other agent has valuation 0 for the first two items and 1 for each of the remaining ones. See the next table.

agent	1	2	3	\dots	n
u	1	n	\cdot	0	\cdot
v	0	1	\cdot	0	\cdot
\cdot	\cdot	\cdot	\cdot	\cdot	\cdot
x	0	0	\cdot	1	\cdot
\cdot	\cdot	\cdot	\cdot	\cdot	\cdot

The allocation in which agents u and v get items 1 and 2, respectively, and each of the remaining agents get one of the items 3 to n is G -EF (recall that $(u, v) \notin E(G)$). On the other hand, no allocation can be PROP as each of the agents u and v need item 2 in order to reach her proportionality threshold for the items allocated to her and her non-neighbors. \square

⁵We note that we have adapted the original definition of graph envy-freeness, which uses undirected graphs in the paper of Chevaleyre, Endriss and Maudet (2017), in order to be as close to the G -PEF fairness notion as possible. Still, G -PEF is much stronger.

6 A General Framework for Defining Fairness Concepts

We now propose an abstract and powerful framework of fairness concepts that captures all existing concepts (previous ones and the ones in the current paper) for indivisible items and seems to be useful in defining new ones. The framework is based on a profile of hypergraphs, with a hypergraph corresponding to each agent. Let $\mathcal{H} = (\mathcal{H}_1, \dots, \mathcal{H}_n)$ be a hypergraph profile where each \mathcal{H}_i is a hypergraph over $[n]$ consisting of hyper-edges that contain node i . Let X be a fairness concept such as EF, PROP, mFS, MFS, etc.

An allocation A satisfies \mathcal{H} -HG- X if for any agent $i \in [n]$ and any hyperedge $e \in \mathcal{H}_i$, there is a re-allocation $A_i^{i,e}$ of the items in $\cup_{j \in e} A_j$ among the agents of e such that $A_i^{i,e} = A_i$, in which the fairness constraints corresponding to fairness concept X for agent i with respect to the agents in e are satisfied. Let us specify the fairness constraints corresponding to fairness concepts EF, PROP, mFS, and MFS for agent i in the allocation $A_i^{i,e}$:

- EF: $v_i(A_i^{i,e}) \geq v_i(A_j^{i,e})$ for every $j \in e$.
- PROP: $v_i(A_i^{i,e}) \geq \frac{1}{|e|} \sum_{j \in e} v_i(A_j^{i,e})$.
- mFS: $v_i(A_i^{i,e}) \geq u_i^{\text{mFS}}(\cup_{j \in e} A_j^{i,e}, |e|)$.
- MFS: $v_i(A_i^{i,e}) \geq u_i^{\text{MFS}}(\cup_{j \in e} A_j^{i,e}, |e|)$.

Observe that, as $v_i(A_i^{i,e}) = v_i(A_i)$, $\sum_{j \in e} v_i(A_j^{i,e}) = \sum_{j \in e} v_i(A_j)$, $u_i^{\text{mFS}}(\cup_{j \in e} A_j^{i,e}, |e|) = u_i^{\text{mFS}}(\cup_{j \in e} A_j, |e|)$, and $u_i^{\text{MFS}}(\cup_{j \in e} A_j^{i,e}, |e|) = u_i^{\text{MFS}}(\cup_{j \in e} A_j, |e|)$, in all cases above besides EF, the re-allocation $A_i^{i,e}$ does not need to be different than A .

Intuitively, hyper-edges of \mathcal{H}_i indicate social constraints that affect the awareness level of agent i about the allocation. A hyper-edge $e \in \mathcal{H}_i$ indicates that agent i knows the items that are allocated to the agents in e besides her, even though it does not provide any information about how they are distributed among these agents. Of course, each agent always knows the items in her bundle.

We can show that several fairness notions from the recent literature as well as all fairness notions defined in the previous sections can be obtained as special cases of this framework. Let us consider some examples.

First, consider the hypergraph profile $\mathcal{H} = (\mathcal{H}_1, \dots, \mathcal{H}_n)$ in which \mathcal{H}_i consists of all pairs of agents that include i . Then, \mathcal{H} -HG-EF is identical to EF. Actually, \mathcal{H} -HG-PROP (and, consequently, \mathcal{H} -HG-mFS) is identical to EF as well since a PROP allocation among two agents is always EF. Notice that, as long as all pairwise hyper-edges are present in all hypergraphs of the profile \mathcal{H} , we cannot get any stronger fairness concept than EF by allowing bigger hyper-edges. \mathcal{H} -HG-MFS is the pairwise-MFS fairness concept from Caragiannis et al. (2016) and requires that the restriction of the allocation to any pair of agents is MFS. We remark that it is not known whether all resource allocation instances admit such allocations or not.

Next, consider the hypergraph profile in which \mathcal{H}_i consists of the single hyper-edge $[n]$ for every $i \in [n]$. In this

case, \mathcal{H} -HG-EF is equivalent to EEF. Using mFS, PROP, or MFS for X instead of EF, \mathcal{H} -HG-X coincides with X.

Let us now redefine the graph versions of fairness notions we studied in Sections 4 and 5. For a given simple graph G , we define the hypergraph profile as follows: \mathcal{H}_i contains the hyperedge $[n] \setminus \text{Neig}_G(i)$ (which consists of i and all her non-neighbors in G) and pair $\{i, j\}$ for each neighbor j of i in G . Then, \mathcal{H} -HG-EF and \mathcal{H} -HG-PROP are identical to G -EEF and G -PEF, respectively.

Finally, given a simple graph G , consider the hypergraph profile $\mathcal{H} = (\mathcal{H}_1, \dots, \mathcal{H}_n)$ in which \mathcal{H}_i contains the pair $\{i, j\}$ for each neighbor j of i in G . Then, \mathcal{H} -HG-EF is the graph envy-freeness notion from Chevaleyre, Endriss and Maudet (2017) while \mathcal{H} -HG-PROP is the (discrete version of) network proportionality that is studied by Abebe, Kleinberg and Parks (2017); Bei, Qiao and Zhang (2017).

7 Future Research Directions

Let us conclude by briefly discussing possible directions for future research. First, there are other epistemic notions of fairness that are worth studying. As an example, let us use the term *optimistic EEF* for an allocation which, for every agent i and every re-allocation of the remaining items, there is another agent j that agent i does not envy. How does optimistic EEF relate to known fairness notions? Clearly, it is weaker than proportionality and it is not hard to see that this property is not weaker than MFS (consider two agents of identical valuations and one item).

In Sections 4 and 5, we have combined only PROP and EF with social graphs; mFS and MFS could also be used. EFX — envy-freeness up to any single item; see Caragiannis et al. (2016) — is another basic fairness concept that could also be used, either together with social graphs or as part of our abstract framework. This could lead to another rich sub-hierarchy of fairness notions but would require the existence of non-EFX resource allocation instances; whether such instances exist or not is a major open problem.

Novel fairness concepts can be defined using our abstract framework. Consider, for instance, \mathcal{H} -HG-PROP in which \mathcal{H}_i consists of all k -sized sets of agents that contain i . This yields EF for $k = 2$, PROP for $k = n$, and new intermediate fairness concepts for the values in between.

Finally, it would be interesting to extend the fairness notions defined here to agents with non-additive valuations or in the completely different settings of (contiguous) cake-cutting or allocations of divisible items.

Acknowledgments

Haris Aziz is supported by a Julius Career Award. Sylvain Bouveret and Jérôme Lang are supported by project ANR-14-CE24-0007-01 “CoCoRiCo-CoDec”. The authors also thank Abdallah Saffidine for insightful discussions.

References

Abebe, R.; Kleinberg, J. M.; and Parkes, D. C. 2017. Fair division via social comparison. In *Proceedings of the 16th Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 281–289.

Bei, X.; Qiao, Y.; and Zhang, S. 2017. Networked fairness in cake cutting. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, 3632–3638.

Bouveret, S., and Lemaître, M. 2016. Characterizing conflicts in fair division of indivisible goods using a scale of criteria. *Autonomous Agents and Multiagent Systems* 30(2):259–290.

Bouveret, S.; Cechlárová, K.; Elkind, E.; Igarashi, A.; and Peters, D. 2017. Fair division of a graph. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, 135–141.

Bouveret, S.; Chevaleyre, Y.; and Maudet, N. 2016. Fair allocation of indivisible goods. In Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D., eds., *Handbook of Computational Social Choice*. Cambridge University Press.

Brams, S. J., and Taylor, A. D. 1996. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press.

Budish, E. 2011. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy* 119(6):1061–1103.

Caragiannis, I.; Kurokawa, D.; Moulin, H.; Procaccia, A. D.; Shah, N.; and Wang, J. 2016. The unreasonable fairness of maximum nash welfare. In *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*, 305–322.

Chen, Y., and Shah, N. 2017. Ignorance is often bliss: Envy with incomplete information. Manuscript.

Chevaleyre, Y.; Dunne, P. E.; Endriss, U.; Lang, J.; Lemaître, M.; Maudet, N.; Padget, J. A.; Phelps, S.; Rodríguez-Aguilar, J. A.; and Sousa, P. 2006. Issues in multiagent resource allocation. *Informatica* 30(1):3–31.

Chevaleyre, Y.; Endriss, U.; and Maudet, N. 2017. Distributed fair allocation of indivisible goods. *Artificial Intelligence* 242:1–22.

Lipton, R. J.; Markakis, E.; Mossel, E.; and Saberi, A. 2004. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Economics and Computation (EC)*, 125–131.

Nguyen, T., and Rothe, J. 2014. Minimizing envy and maximizing average nash social welfare in the allocation of indivisible goods. *Discrete Applied Mathematics* 179:54–68.

Plaut, B., and Roughgarden, T. 2017. Almost envy-freeness with general valuations. *CoRR* abs/1707.04769.

Procaccia, A. D., and Wang, J. 2014. Fair enough: guaranteeing approximate maximin shares. In *Proceedings of the 15th ACM Conference on Economics and Computation (EC)*, 675–692.

Procaccia, A. D. 2016. Cake cutting algorithms. In Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D., eds., *Handbook of Computational Social Choice*. Cambridge University Press.

Robertson, J. M., and Webb, W. A. 1998. *Cake Cutting Algorithms: Be Fair If You Can*. A. K. Peters.

Steinhaus, H. 1948. The problem of fair division. *Econometrica* 16(1):101–104.