



Fair Division via Social Comparison

Citation

Abebe, Rediet, Jon Kleinberg, and David Parkes. 2017. Fair Division via Social Comparison. AAMAS '17 Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, São Paulo, Brazil, May 08- 12, 2017, 281-289.

Published version

<https://dl-acm-org.ezp-prod1.hul.harvard.edu/citation.cfm?id=3091171>

Link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:41593087>

Terms of use

Metadata Only

Accessibility

<https://accessibility.huit.harvard.edu/digital-accessibility-policy>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#)

Fair Division via Social Comparison

Rediet Abebe
Cornell University
Ithaca, NY 14853
red@cs.cornell.edu

Jon Kleinberg
Cornell University
Ithaca, NY 14853
kleinber@cs.cornell.edu

David C. Parkes
Harvard University
Cambridge, MA 02138
parkes@eecs.harvard.edu

ABSTRACT

We study cake cutting on a graph, where agents can only evaluate their shares relative to their neighbors. This is an extension of the classical problem of fair division to incorporate the notion of social comparison from the social sciences. We say an allocation is *locally envy-free* if no agent envies a neighbor’s allocation, and *locally proportional* if each agent values its own allocation as much as the average value of its neighbors’ allocations. We generalize the classical “Cut and Choose” protocol for two agents to this setting, by fully characterizing the set of graphs for which an oblivious *single-cutter protocol* can give locally envy-free (thus also locally-proportional) allocations. We study the *price of envy-freeness*, which compares the total value of an optimal allocation with that of an optimal, locally envy-free allocation. Surprisingly, a lower bound of $\Omega(\sqrt{n})$ on the price of envy-freeness for global allocations also holds for local envy-freeness in any connected graph, so sparse graphs do not provide more flexibility asymptotically with respect to the quality of envy-free allocations.

Keywords

Social Choice; Fair Division; Graph-Theoretic Methods

1. INTRODUCTION

The fair allocation of resources is a fundamental problem for interacting collections of agents. A central issue in fair allocation is the process by which each agent compares its allotment to those of others’. While theoretical models have tended to focus on global comparisons — in which an agent makes comparisons to the full population — a rich line of empirical work with its origins in the social sciences has suggested that in practice, individuals often focus their comparisons on their social network neighbors. This literature, known as *social comparison theory*, dates back to work of Festinger [16], and has been explored extensively by economists and sociologists since; for example, see Akerlof [2] and Burt [10]. The primary argument is that in many contexts, an individual’s view of their subjective well-being is based on a comparison with peers, defined through an underlying social network structure, rather than through com-

parison with the overall population [21]. For instance, when people take part in a labor market, they can in principle apply to any available job, and we therefore have a global matching market. But, social comparison theory suggests that in many contexts, they will evaluate their job outcome by how it turned out in relation to their social peers [10]. This distinction between the global process of allocation and the local definition of fairness is one of the underpinnings of social comparison.

In this work, we find that the perspective of social comparison theory motivates a rich set of novel theoretical questions in classical resource allocation problems. In particular, we apply this theory to the canonical *cake cutting* problem, which refers to the challenge of allocating a single divisible, continuous, good in a fair and efficient manner. The “cake” is intended to stand for a good over which different agents have different preferences for difference pieces. This problem has a wide range of applications including international border settlements, divorce and inheritance settlements, and allocating shared computational resources.

The cake is represented using the $[0, 1]$ interval. Agent preferences are modeled through functions that map subintervals to real numbers according to the value the agent assigns to that piece. We normalize these valuations so that each agent’s value for the whole cake is 1. The entire cake is to be allocated, but agents need not receive a single, contiguous interval (and valuations are additive across pieces).

Following our goal of understanding the properties of local comparisons to network neighbors, we study cake cutting in a setting where there is an underlying network on the agents, and fairness considerations are defined locally to an agent’s neighbors’ in the network. Note here that the social network is not imposing a constraint on the allocation procedure itself, but rather how the agents evaluate their allocated pieces. Given a directed graph G and a cake to be allocated, we define a *locally proportional allocation* to be one where each agent values her allocation at least as much as the average value she would obtain from her neighbors’ allocations in G .¹ We define a *locally envy-free allocation* to be one where no agent envies the allocation of any neighbor in G . Note, the typical notions of fairness are attained by replacing the neighborhood of an agent with the set of all other agents.

Appears in: *Proc. of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2017)*, S. Das, E. Durfee, K. Larson, M. Winikoff (eds.), May 8–12, 2017, São Paulo, Brazil.
Copyright © 2017, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

¹For example, if someone asks, “Is my salary at least as high as the average salary of my co-workers?”, they’re asking a local proportionality question. This is more typical than the corresponding global one: “Is my salary at least as high as the average salary of all humans in the world?”

As in the global case, it is straightforward to see that a locally envy-free allocation is also locally proportional. It is also clear that if H is a subgraph of G on the same node set, then a locally envy-free allocation for G must also be locally envy-free for H . In particular, global envy-freeness implies local envy-freeness. On the other hand, we show that local and global proportionality are incomparable in general.

To avoid direct revelation of valuation functions, as is common in the cake cutting literature, we use the Robertson-Webb model [25], where protocols that interact with agents through two types of queries: $\text{cut}(x, y, \alpha)$ and $\text{eval}(x, y)$, for $x, y, \alpha \in [0, 1]$. A *cut* query asks an agent for a cut-point x such that the subinterval $[x, y]$ has value α to the agent. The *eval* query asks for the agent’s value for the subinterval $[x, y]$. These queries provide information about valuations in order to compute an allocation. The *query complexity* of a cake cutting problem is the worst case number of queries required in order to output an allocation satisfying the desired fairness criteria.

Network topology plays a crucial role in our results. We look for non-trivial classes of graphs for which we can give efficient protocols for (locally) fair allocations. We start from the structure of the classical “cut-and-choose” solution for two agents: one agent divides the cake and the other selects a piece. Our first result is that we fully characterize the family of graphs for which a locally envy-free (and thus also locally proportional) allocation can be produced by a protocol in which a single designated agent performs all the cuts at the outset, based only on her valuation function. In contrast, cut-and-choose does not provide any guarantees for global allocations with more than two agents.

We also consider the effect of fairness on welfare. The *price of fairness* [11] is defined as the worst case ratio over all inputs between the social welfare of the optimal allocation (the allocation maximizing the sum of agent valuations) and the social welfare of the optimal fair allocation. We refer to this ratio in the case of envy-free allocations as the *price of (local) envy-freeness*, and in the case of proportional allocations as the *price of (local) proportionality*. Caragiannis et al. [11] give an $\Omega(\sqrt{n})$ lower bound for the price of global envy-freeness and proportionality, a matching upper bound of $O(\sqrt{n})$ for the price of global proportionality, and a loose upper bound of $n - 1/2$ for the price of global envy-freeness.

Our second, main result is to show that the lower bound of $\Omega(\sqrt{n})$ also holds on the price of local envy-freeness in any connected, undirected graph G . This is surprising since one might expect sparse graphs to provide more flexibility with respect to the quality of envy-free allocations. The known upper bound of $n - 1/2$ for the price of global envy-freeness serves as a loose upper bound for the price of local envy-freeness (and also local proportionality). Whether there is a similar lower bound for local proportionality and tighter upper bounds are left as open questions for future work.

Related Work.

A proportional cake cutting protocol for any number of agents was found by Evan and Paz [15]. It was later shown that one cannot do better than their $\Theta(n \log n)$ query complexity [20]. Progress for envy-free allocations has been slower. The cut-and-choose solution for two agents can be dated back to the Book of Genesis. For three agents, bounded protocols were given independently by Selfridge (1960) and Conway (1996) [8]. In 1995, Brams and Taylor [7] gave an envy-

free, but potentially unbounded, protocol for any number of agents. In a breakthrough result, Aziz and Mackenzie [3] provided a bounded protocol for $n = 4$, and shortly afterward a bounded protocol for any n [4]. Although bounded, the protocol has very high multiple-exponential query complexity, while the best lower bound is $\Omega(n^2)$ [22].

Due to the difficult nature of the envy-free cake cutting problem, researchers have imposed restrictions on different aspects in order to give useful protocols and gain insight. A few examples are: restricting valuation functions to be only piecewise constant or uniform [19], relaxing envy-freeness to approximate envy-freeness [24], considering partial allocations that simultaneously satisfy envy-freeness and proportionality [12], and limiting allocations to be contiguous pieces [28]. A summary of several classical results is given by see Brams and Taylor [8]. Procaccia [23, 24] gives a recent survey from a computer science perspective. Working outside the Robertson-Webb model, Chen et al. [12] provide a strategy-proof mechanism for fair allocations for piecewise uniform valuation functions.

We are not aware of existing work on the cake cutting problem in which fairness is determined via comparisons defined by an underlying graph. Most related is the work of Chevaleyre et al. [13], who study the allocation of indivisible goods on a network. In their model, an agent only envies allocations of those in its neighborhood. In addition, the network constrains the agents with which one can negotiate and thus the allowable interactions in the division protocol. In a somewhat different direction, there has been research in which valuation functions are generalized to include externalities in the sense that agents can derive valuations from other agent’s allocations as well [9]. Finally, at a more general level, the line of work on graphical games has studied other contexts in which games among multiple agents are structured so that each agent’s payoffs depend only the interactions with network neighbors [17].

2. RELATING GLOBAL AND LOCAL PROPERTIES

Let $N = \{1, 2, \dots, n\}$ denote the set of agents. The cake is represented using the interval $[0, 1]$ and a *piece of cake* is a finite union of non-overlapping (interior disjoint) subintervals of $[0, 1]$. Allocated pieces are a finite union of subintervals. Each agent i has a valuation function V_i that maps subintervals to values in \mathbb{R} . Given subinterval $[x, y] \subseteq [0, 1]$, we write $V_i(x, y)$ instead of $V_i([x, y])$ for simplicity. We assume that valuation functions are additive, non-atomic, and non-negative. Non-atomicity gives us $V_i(x, x) = 0$ for all $x \in [0, 1]$, so we can ignore boundaries when defining cut-points. We normalize valuations so that $V_i(0, 1) = 1$ for each agent i .

Definition 1 (Allocation). *An allocation is a partition of the $[0, 1]$ interval into n pieces $\{A_1, A_2, \dots, A_n\}$ such that $\cup_i A_i = [0, 1]$ and the pieces are pairwise disjoint. Each agent i is assigned the corresponding piece A_i .*

As is standard, this ensures that the entire cake is allocated. If we remove this constraint, then we can have trivial solutions that satisfy fairness, such as assigning each agent nothing in the case of envy-freeness. This assumption is a natural one to make since the valuation functions are non-negative and additive and thus satisfy free-disposal.

The Robertson-Webb query model for *cake cutting protocols* is defined with the following two types of queries:

- $\text{eval}_i(x, y)$; this asks agent i for the valuation $V_i(x, y)$.
- $\text{cut}_i(x, y, \alpha)$: given $y, \alpha \in [0, 1]$, this asks agent i to pick $x \in [0, 1]$ such that $V_i(x, y) = \alpha$.

These queries are used to gather information regarding the valuations of the agents and need not directly determine an allocation. Rather, a cake cutting protocol can use other steps for determining allocations.

Definition 2 (Query complexity). *The query complexity of a cake cutting protocol is the worst case number of queries that the protocol requires to output an allocation over all possible valuation functions.*

The query complexity of a cake cutting problem is the minimum query complexity over all known protocols for computing the desired allocation.

2.1 Global and Local Fairness

Given a set of agents and allocation $\mathcal{A} = (A_1, A_2, \dots, A_n)$, we formally define two global fairness criteria:

Definition 3 (Proportional, Envy-free). *An allocation \mathcal{A} is proportional if $V_i(A_i) \geq 1/n$, for all $i \in N$, and is envy-free if $V_i(A_i) \geq V_i(A_j)$, for all $i, j \in N$.*

Suppose we are given a directed graph $G = (V, E)$, where the nodes correspond to agents and edges signify relations between the agents. In particular, we assume that given a directed edge (i, j) , agent i can view agent j 's allocation. Agent i 's *neighborhood* is the set of all nodes to which it has directed edges (i, j) , and we denote this set of nodes by N_i . We define i 's *degree* to be $d_i = |N_i|$. We define local analogues for fairness concepts:

Definition 4 (Local proportional, local envy free). *Given a graph G , an allocation \mathcal{A} is locally proportional if $V_i(A_i) \geq \frac{\sum_{j \in N_i} V_i(A_j)}{|N_i|}$ for all i and $j \in N_i$ and locally envy-free if $V_i(A_i) \geq V_i(A_j)$.*

In a locally proportional allocation, each agent assigns as much value to her allocation as the average value she has for a neighbors' allocation. In a locally envy-free allocation, each agent values her allocation at least as much as her neighbors' allocation.

When $G = K_n$, the complete graph on n vertices, these local fairness definitions coincide with their global analogues. Whereas, if $G = I_n$, the empty graph on n nodes, then any allocation is trivially locally envy-free. So, the graph topology plays a significant role in computing locally fair allocations.

Lemma 5. *A locally envy-free allocation \mathcal{A} on some graph G is also locally envy-free on all subgraphs $G' \subseteq G$.*

PROOF. We want to show that given a node u and $v \in N_u$, u does not envy v 's allocation in G' . This follows from the fact that \mathcal{A} is a locally envy-free allocation on G , and if (u, v) is an edge in G' , then it is also an edge in G .

One consequence of this lemma is that local envy-freeness is implied by global envy-freeness. Since globally envy-free allocations exist for all sets of agent valuations [1], a locally envy-free allocation exists for every graph G and every set of agent valuations.

Lemma 6. *If an allocation \mathcal{A} is locally envy-free on a graph G , then it is also locally proportional on the same graph.*

PROOF. If an allocation $\mathcal{A} = (A_1, A_2, \dots, A_n)$ is locally envy-free, then for any $i \in V$, $V_i(A_i) \geq V_i(A_j), \forall j \in N_i$. Therefore, $V_i(A_i) \geq (\sum_{j \in N_i} V_i(A_j))/|N_i|$.

Therefore, locally proportional allocations also exist. By considering $G = K_n$, we also recover that global envy-freeness implies global proportionality. While global envy-freeness implies local envy-freeness, global proportionality does not necessarily imply local proportionality, or vice versa, the former of which violates intuition. We provide a counter example.

Example 7. *Let $n = 4$ and $G = C_4$, the cycle graph on 4 nodes, where the nodes are labeled clockwise. Assume agents 2, 3, and 4 have the uniform valuation function $V_i(x, y) = |y - x|$ for any subinterval $(x, y) \subseteq [0, 1]$. Let agent 1 have the piecewise uniform valuation function where $V_1(0, 1/4) = 1/2$ and $V_1(3/4, 1) = 1/2$, and no value for the remaining subinterval. It is easy to verify that the following allocation is locally proportional on K_4 ,*

$$\mathcal{A} = ([0, 1/8], [1/8, 3/8], [3/8, 5/8], [5/8, 1]).$$

In particular, $V_i(A_i) = 1/4$ for $i \in \{1, 2, 3\}$ and $V_4(A_4) = 3/8$. This allocation is however not locally proportional on C_4 since $V_1(A_1 \cup A_2 \cup A_4) = 1$, but $V_1(0, 1/8) = 1/4 < 1/3$. It is also not locally envy-free since $V_1(A_4) > V_1(A_1)$.

We prove a stronger result regarding any pair of distinct graphs. Note by $N_i(H)$ we mean agent i 's neighborhood set in graph H .

Theorem 8. *Given any pair of distinct, connected graphs G, H on the same set of nodes, there exists a valuation profile of the agents and an allocation \mathcal{A} such that \mathcal{A} is locally proportional on G but not on H .*

PROOF. First, consider the case where H is a strict subgraph of G . Pick a node i such that $|N_i(H)| < |N_i(G)|$. Let $N_i(G) = \{i_1, i_2, \dots, i_k\}$ and $N_i(H) = \{i_1, i_2, \dots, i_\ell\}$ for some $\ell < k$. Assume that all other nodes besides i have a uniform valuation function over the entire cake. Then, the allocation $A_j = ((j-1)/n, j/n)$ is locally proportional from the perspective of every other agent $j \in N$ on both H and G . Now, define i 's valuation function to be the piecewise uniform valuation function where, $V_i((i-1)/n, i/n) = 1/(|N_i(G)|+1)$ and $V_i(A_{i_1} \cup A_{i_2} \dots \cup A_{i_\ell}) = 1 - 1/(|N_i(G)|+1)$. Agent i 's valuation for the allocations of agents in the set $\{i_{\ell+1}, i_{\ell+2}, \dots, i_k\}$ as well as $V(G) \setminus N_i(G)$ is 0. This allocation \mathcal{A} is therefore locally proportional on G . For \mathcal{A} to be locally proportional on H , we need $V_i(A_i) \geq 1/|N_i(H)|$. However, since $|N_i(H)| < |N_i(G)|$, and $V_i(A_i \cup A_{i_1} \cup A_{i_2} \dots \cup A_{i_\ell}) = 1$, we only have that $V_i(A_i) < 1/(|N_i(H)|+1)$.

Now, suppose $H \not\subseteq G$. Then, there exists an edge (i, j) in the edge-set of H that is not in the edge-set of G . Assume that all nodes $k \neq i$ have a uniform valuation over the entire cake. Suppose further that we have the allocation where each k is assigned the piece $A_k = ((k-1)/n, k/n)$. As above, this allocation is locally proportional from the perspective of each agent k on both G and H . Define i 's valuation function to be $V_i((j-1)/n, j/n) = 1$ and 0 on the remainder of the cake. Then, $V_i(A_i) = 0$ and $V_i(A_k) = 0$ for all $k \neq j$. Since $j \notin N_i(G)$, this allocation is locally proportional on G . However, it is not locally proportional on H since $V_i(\cup_{\ell \in N_i} A_\ell) = 1$, but $V_i(A_i) = 0$.




Figure 1: Relationship Between Fairness Concepts

3. ENVY-FREE NETWORK ALLOCATIONS

In this section, we consider the question of finding efficient protocols for computing locally envy-free allocations. We assume that graphs G are directed, unless specified otherwise. When we mean the *component* of a directed graph, we will instead take the graph obtained by replacing each directed edge with an undirected one, and a component in the directed graph is the corresponding subgraph to the connected component in the undirected analogue. We will use *strongly connected component* when we mean to take directed reachability into account.

Lemma 9. *Suppose we have a bounded protocol for computing locally envy-free allocations on G . The same protocol can be used to compute locally envy-free allocations on the following two classes of graphs: (i) $H = G \cup G'$ where G and G' are disjoint components, and (ii) when H is a graph with a directed cut such that every edge across the cut goes from a node in G to a node in $H \setminus G = G'$.*

PROOF. In both instances, we simply apply the protocol on G and allocate agents in G' the empty allocation. This is a locally envy-free allocation on H , since no two agents in G envy one another by the assumption on the protocol, and no agent envies the allocation of another agent in G' .

A few consequences of this lemma are that: given a graph G with more than one connected component, we can reduce the search for a bounded protocol on G to any one of the connected components. Furthermore, if G is a *directed acyclic graph* (DAG), then there exists at least one node with no incoming edges. Therefore, the allocation where such a node gets the entire cake—or where it is divided among a set of such nodes—is locally envy-free.

3.1 Directed Acyclic Graphs and Their Cones

We consider a conceptually useful class of graphs for which we can give a protocol with query complexity of $O(n^2)$.

Definition 1. *Given a graph $G = (V, E)$, we say that G' is a cone of G if it is the join of G and a single node c , which we call the apex. That is, G' has node set $V \cup \{c\}$, and edge set consisting of the edges of G , together with undirected edges (u, c) for all $u \in V$. We denote the cone G' of G by $G \star c$.*

We consider cones of DAGs. These are the class of graphs where there is a single node c that lies on all cycles. This class of graphs include many interesting classes including all graphs, where each node is included in at most one cycle independent of the number of agents, and we will use this to present a procedure for extending partial envy-free allocations to full ones, generalizing a result by [3].

We now show how to compute a locally envy-free allocation on any graph that is the cone of a DAG.




Figure 2: Cone of a Directed Acyclic Graph

Protocol 1: Cone of DAGs

- 1: Agent c cuts the cake into n pieces that she values equally.
 - 2: Topologically sort and label the nodes $N \setminus \{c\}$ such that for every edge (i, j) , $i \leq j$.
 - 3: Nodes $N \setminus \{c\}$ pick a piece they prefer most in increasing order of their index.
 - 4: Agent c takes the remaining piece.
-

Theorem 10. *Given a graph G that is a cone of a DAG, Protocol 1 computes a locally envy-free allocation on G using a bounded number of queries.*

PROOF. We first show that the allocation is locally envy-free. First, there is no envy between agent c and any other agent since agent c cuts the cake into n pieces she values equally. Therefore, her valuation for all the allocated pieces is $1/n$. Each of the other agents picks a piece before c , and so are able to pick a piece that they value at least as much as the remaining piece that agent c is assigned. Finally, given any directed edge (i, j) such that $i, j \neq c$, note that $V_i(A_i) \geq V_i(A_j)$, since if such an edge exists, then $i < j$ and thus i selects a piece before j .

To count the number of queries, the first step requires $n-1$ cut queries by agent c . Then, each agent i must perform $n-i+2$ eval queries to determine the piece for which they have the highest value. Therefore, the protocol above uses $(n^2 + 3n - 4)/2$ queries.

The importance of cones of DAGs can be seen in the following result, which shows that they emerge naturally as the characterization of graphs on which a particular fundamental kind of protocol succeeds.

Definition 11. *An oblivious single-cutter protocol is one in which a single agent i first divides up the cake into a set of pieces P_1, P_2, \dots, P_t (potentially $t > n$), and then all remaining operations consist of agents choosing from among these pieces.*

The classical Cut-and-Choose protocol is an oblivious single-cutter protocol that works for all sets of valuation functions on the complete, two-node graph K_2 . Protocol 1 is an oblivious single-cutter protocol which works for any graph that is the cone of a DAG. We show that subgraphs of cones of DAGs are in fact precisely the graphs on which oblivious single-cutter protocols are guaranteed to produce a locally envy-free allocation.

Theorem 12. *If G is a graph for which an oblivious single-cutter protocol produces a locally envy-free allocation for all sets of valuation functions, then G is a subgraph of the cone of a DAG.*

PROOF. Suppose, by way of contradiction, that G is not a subgraph of the cone of a DAG, but that there is an oblivious single-cutter protocol on G , in which a node i starts by dividing the cake into pieces using only knowledge of her own valuation function. Since G is not a subgraph of the cone of a DAG, the graph $G_i = G \setminus \{i\}$ is not acyclic, so there is a cycle $C = (c_1, c_2, \dots, c_m, c_1)$ in G_v .

Let i have a valuation function such that she produces a partition of the cake into pieces P_1, P_2, \dots, P_t . Because agent i produces these pieces without knowledge of the valuation functions of the other agents, we can imagine that we adversarially choose the valuations of the other agents after these pieces have been produced. In particular, consider the valuation functions in which each node c_j on the cycle C values piece P_r (for $r < t$) at $2 \cdot 3^{-r}$, and the last piece P_t with the remaining value. These valuations have the property that for each r , the piece P_r is more valuable than the union of all pieces $P_{r+1} \cup P_{r+2} \cup \dots \cup P_t$.

After the protocol is run, each agent c_j on C will get a subset of the pieces produced by i . Let s be the minimum index of any piece allocated to an agent c_j on C . Then, the agent c_{j-1} who has a directed edge to c_j will have a union of pieces that she values less than she would value P_s , and hence envies c_j . This contradicts the assumption that the protocol produces a locally envy-free allocation on G .

We highlight an important connection between computing locally envy-free allocations on graphs and what is known in the literature as *irrevocable advantage*. Given a partial allocation, an agent i is said to have irrevocable advantage over agent j (or *dominate* j) if agent i remains unenvied of agent j 's allocation even if the entire remaining piece of the cake (the *residue*) is added to agent j 's allocation.

With the additional guarantee that each agent dominates some number of other agents, this concept is often used to extend partial globally envy-free allocations to complete ones. For instance, it is a key concept in the Aziz-Mackenzie protocol for K_4 . Their protocol can be decomposed to three subprotocols: Core, Permutation, and Post-Double Domination Protocols, in order. The *Core Protocol* computes partial envy-free allocation where each agent dominates at least two other agents, while the *Post Double Domination Protocol* extends this to a complete allocation. We will use Protocol 1 to show that given a partial envy-free allocation on K_n where each agent dominates at least $n - 2$ other agents, we can extend the allocation to a complete one, thereby generalizing the Post Double Domination Protocol [3] for any n . This is presented in Appendix A.

4. PRICE OF FAIRNESS

Finally, we consider the efficiency of allocation from the perspective of local fairness. We follow the approach introduced by Caragiannis *et al.* [11] of studying the price of envy-freeness, and for this we begin with the following definitions. Recall that for an allocation \mathcal{A} into pieces $\{A_1, A_2, \dots, A_n\}$ for the n agents, we use $V_i(A_i)$ to denote agent i 's valuation for its piece.

Definition 13 (Optimality). *An allocation \mathcal{A} , is said to be optimal if $\sum_i V_i(\mathcal{A}_i) \geq \sum_i V_i(\mathcal{B}_i)$ for any allocation \mathcal{B} . We denote this optimal allocation by \mathcal{A}^* .*

We define the *optimal locally envy-free* (resp. *optimal locally proportional*) allocations, denoted by $\mathcal{A}^{\text{LEF}^*}$ (resp.

$\mathcal{A}^{\text{LP}^*}$), analogously by imposing the constraint that \mathcal{A} and \mathcal{B} be locally envy-free (resp. locally proportional) and maximizing sum of the values across all agents.

Definition 14 (Price of Local Envy-Freeness, Proportionality). *Given an instance of a cake cutting problem on a graph G , the price of local envy-freeness is the ratio,*

$$\frac{\sum_i V_i(\mathcal{A}_i^*)}{\sum_i V_i(\mathcal{A}_i^{\text{LEF}^*})},$$

where the sum is over all agents $i \in N$. We likewise define the price of proportionality by taking the denominator to be $\sum_i V_i(\mathcal{A}_i^{\text{LP}^*})$.

We are measuring the degradation in efficiency when considering allocations that maximize the welfare in both instances under the given constraints. To quantify the loss of efficiency, we are interested in giving a tight lower and upper bound. More specifically, given a graph G and a fairness concept in consideration, say local envy-freeness, we seek to find an input (i.e., a valuation profile) for which the price of local envy-freeness is high. This corresponds to a lower bound on the price of fairness. On the other hand, the upper bound will be given via an argument that shows, for any valuation profile, the price of fairness cannot exceed that stated.

The main result on global envy-freeness, due to Caragiannis *et al.* [11] is an $\Omega(\sqrt{n})$ lower bound on the price of (global) fairness: there exist valuation functions for which the ratio is $\Omega(\sqrt{n})$. (Very little is known about the upper bound for the price of envy-freeness: an upper bound of n is immediate, and the best known upper bound is $n - 1/2$ [11].)

These existing results are for the standard model in which each agent can envy every other agent. Using our graph-theoretic formulation, we can study the price of local fairness. As defined in Section 2, this is the ratio of the total welfare of the optimal allocation to the maximum total welfare of any allocation that is locally envy-free.

The numerator of this ratio—based on the optimal allocation—is independent of G , while the denominator is a maximum over a set of allocations that is constrained by G . Now, if we imagine reducing the set of edges in G , the set of allocations eligible for the maximum in the denominator becomes less constrained; consequently, we would expect that the price of fairness may become significantly smaller as G becomes sparser. Is this in fact the case? We show that it is not. Our main result is that the lower bound for global envy-freeness also applies to local envy-freeness on any connected undirected graph.

Theorem 15. *For any connected undirected graph G , there exist valuation functions for which the price of local envy-freeness on G is $\Omega(\sqrt{n})$.*

To prove this theorem, we start by adapting a set of valuation functions that Caragiannis *et al.* [11] used in their lower bound for global envy-freeness. To argue about their effect on allocations in an arbitrary graph G , we need to reason about the paths connecting agents in G to others with different valuation functions. This, in turn, requires a delicate graph-theoretic definition and argument: we introduce a structure that we term a (k, ϵ) -linked partition; we show that if G contains this structure, then we can carry out the

lower bound argument in G ; and finally we show that every connected undirected graph contains a (k, ε) -linked partition.

Definition 2. For a connected graph $G = (V, E)$, a natural number $k \geq 1$, and a real number $0 < \varepsilon \leq 1$, we define a (k, ε) -linked partition as follows. It consists of a set $L \subseteq V$ of size k , and a partition of $S = V - L$ into sets $\{S_i : i \in L\}$ each of size at least $(\varepsilon n/k) - 1$, such that for each $j \in S_i$, there is an i - j path in $S \cup \{i\}$. That is, each $j \in S_i$ can reach $i \in L$ without passing through any other nodes of L .

We next show that if a connected undirected graph G has such a structure, with appropriate values of k and ε , then we obtain a lower bound on the price of local envy-freeness on G .

Lemma 16. If a connected undirected graph G has a (k, ε) -linked partition with $k = \lfloor \sqrt{n} \rfloor$ and ε , a constant, then there exist valuation functions on the nodes of G for which the price of local envy-freeness is $\Omega(\sqrt{n})$.

PROOF. Suppose G has a (k, ε) -linked partition consisting of L and $\{S_i : i \in L\}$, where $k = \sqrt{n}$. Thus, each S_i has size at least $\varepsilon\sqrt{n} - 1$. (We will assume for the sake of exposition that \sqrt{n} is an integer, although it is straightforward to slightly modify the argument if it is not.)

We now use a valuation function adapted from the construction of Caragiannis et al [11], who considered the price of global envy-freeness. We partition the full resource to be allocated, the interval $[0, 1]$, into \sqrt{n} disjoint intervals $I_1, \dots, I_{\sqrt{n}}$. We will give each $i \in L$ a valuation v_i that places all value on distinct interval I_i , and each $j \in S$ a valuation that is uniform on $[0, 1]$. The optimal allocation for this set of valuations has total welfare of \sqrt{n} , which is achieved by giving each $i \in L$ the entire interval where it places value.

Now let us consider any envy-free allocation $\mathcal{A} = \{A_i : i \in V\}$. Let μ_i be a real number denoting the Lebesgue measure of the set A_i assigned to node i . If $j \in S$, then j 's valuation for its set, $v_j(A_j)$ is equal to μ_j . If $i \in L$, then i 's valuation $v_i(A_i)$ is \sqrt{n} times the measure of $A_i \cap I_i$; hence $v_i(A_i) \leq \mu_i \sqrt{n}$.

For each $i \in L$, each $j \in S_i$ has a path P_j to i entirely through nodes of S ; let the nodes on this path, beginning at i , be $P_j = i, j_1, j_2, \dots, j_d = j$. The immediate neighbor j_1 of i on P_j must satisfy $\mu_{j_1} \geq \mu_i$, since j_1 is in S and hence has a uniform valuation on intervals. For each successive j_t on P , we must have $\mu_{j_t} = \mu_{j_{t-1}}$, since j_t and j_{t-1} have the same valuation on all sets, and the allocation is locally envy-free. Thus, by induction we have $\mu_{j_t} \geq \mu_i$ for all t , and hence $\mu_j \geq \mu_i$.

We can now derive a set of inequalities that establishes the lower bound. First we have,

$$\sum_{j \in S_i} v_j(A_j) = \sum_{j \in S_i} \mu_j \geq \sum_{j \in S_i} \mu_i \geq \mu_i(\varepsilon\sqrt{n} - 1).$$

Let us assume n is large enough that $\varepsilon\sqrt{n} - 1 \geq \varepsilon\sqrt{n}/2$, so we have,

$$\sum_{j \in S_i} v_j(A_j) \geq \mu_i(\varepsilon\sqrt{n}/2).$$

Since $v_i(A_i) \leq \mu_i \sqrt{n}$ for $i \in L$, we have,

$$\sum_{j \in S_i} v_j(A_j) \geq \varepsilon v_i(A_i)/2. \quad (1)$$

Thus, the total welfare of the allocation is,

$$\begin{aligned} \sum_{h \in V} v_h(A_h) &= \sum_{i \in L} v_i(A_i) + \sum_{j \in S} v_j(A_j) \\ &= \sum_{i \in L} [v_i(A_i) + \sum_{j \in S_i} v_j(A_j)] \\ &\leq \sum_{i \in L} [(2\varepsilon^{-1} + 1) \sum_{j \in S_i} v_j(A_j)] \\ &= (2\varepsilon^{-1} + 1) \sum_{i \in L} \sum_{j \in S_i} v_j(A_j) \\ &= (2\varepsilon^{-1} + 1) \sum_{j \in S} v_j(A_j) \leq (2\varepsilon^{-1} + 1), \end{aligned}$$

where the first inequality is by (1) and the second since

$$\sum_{j \in S} v_j(A_j) = \sum_{j \in S} \mu_j \leq 1,$$

because all agents in S get disjoint intervals. Since $(2\varepsilon^{-1} + 1)$ is a constant, while the optimal allocation has total welfare \sqrt{n} , this implies an $\Omega(\sqrt{n})$ lower bound on the price of envy-freeness on G .

Finally, we establish that every connected undirected graph G has a (k, ε) -linked partition for appropriate values of k and ε . We begin by showing that it is enough to find a structure satisfying a slightly more relaxed definition, in which the set L can have more than k elements, and we do not need to include all the nodes of G . Specifically, we have the following definition:

Definition 3. For a connected graph $G = (V, E)$, a natural number $k \geq 1$, and a real number $0 < \varepsilon \leq 1$, we define a (k, ε) -linked subpartition as follows. It consists of a set $L \subseteq V$ of size $\ell \geq k$, together with disjoint subsets $S_1, S_2, \dots, S_\ell \subseteq S = V - L$, each of size of size at least $(\varepsilon n/k) - 1$, such that for each $j \in S_i$, there is an i - j path in $S \cup \{i\}$.

The following lemma says that it is sufficient to find a (k, ε) -linked subpartition.

Lemma 17. If a connected undirected graph G contains a (k, ε) -linked subpartition, then it contains a (k, ε) -linked partition.

PROOF. We start with a (k, ε) -linked subpartition of G , with disjoint sets L of size $\ell \geq k$, and $S_1, S_2, \dots, S_\ell \subseteq S = V - L$. First, for every node $v \notin L \cup S$, we assign it to a subset S_i as follows: we find the shortest path from v to any node in L ; suppose it is to $i \in L$. We add v to S_i . Note that this preserves the property that all S_i are disjoint, and v has a path to i that does not meet any other node of L , since if $h \in L$ were to lie on this path, it would be closer to v than i is.

At this point, every node of G belongs to $L \cup S$. We now must remove nodes from L to reduce its size to exactly k while preserving the properties of a (k, ε) -linked partition. To do this, we choose a node $i \in L$ arbitrary, remove i from L , and remove the set S_i from the collection of subsets. We then assign each node in $S_i \cup \{i\}$ to an existing subset S_h exactly as in the previous paragraph. After this process, the size of L has been reduced by 1, and we still have a partition of $V - L$ into subsets S_i with the desired properties. Continuing in this way, we can reduce the size of L to exactly k , at which point we have a (k, ε) -linked partition.

Finally, we prove the following graph-theoretic result, which together with Lemma 16 establishes Theorem 15.

Theorem 18. *For every $k \geq 2$ and with $\varepsilon = 1/2$, every connected undirected graph has a (k, ε) -linked subpartition.*

PROOF. It is enough to find the required structure on a spanning tree T of G , since if the paths required by the definition exist in T , then they also exist in G . Thus, it is sufficient to prove the result for an arbitrary tree T .

We root T at an arbitrary node, and let X be the set of leaves of T . If $|X| \geq k$, then we can choose any k leaves of T and partition the remaining nodes of T arbitrarily into sets of size $(n - k)/k$ to satisfy the definition. Otherwise, $|X| < k$. In this case, we begin by including all nodes of X in L .

Now, we process the nodes of T , working upward from the leaves, so that when we get to a node v in T , we have already processed all descendants of v . Each node is processed once, and at that point we decide whether to add it to L , and if not which set S_i to place it in, given the current set L . For a node v , we say that w is *downward-reachable* from v if w is a descendent of v , and if the v - w path in T does not contain any internal nodes belonging to L .

When we process a node v , we do one of two things:

- (i) We label v with the name of a node in L that is downward-reachable from v ; or
- (ii) We place v in L .

Let $b = (\varepsilon n/k) - 1$. We perform action (i) if there is any $w \in L$ that is downward-reachable from v , such that there are not yet b nodes labeled with w . In this case, we label v arbitrarily with one such w . Since v and all its descendants are now processed, v will continue to have a path to w that does not pass through any other nodes of L . Otherwise suppose there is no such w ; that is, all $w \in L$ that are downward-reachable from v have b nodes labeled with w . In this case, we perform action (ii). Note that at this point, every $w \in L$ that is a descendent of v has a set S_w of exactly b nodes, and these nodes can all reach w without passing through any other node of L .

Our procedure comes to an end when we process the root node v^* . There are three cases to consider, the first two of which are straightforward.

First, if we place v^* into L , then $T - v^*$ is partitioned into L and sets $\{S_w : w \in L\}$ such that $|S_w| = b$ for each w . Thus, if we remove v^* from L , we have a (k, ε) -linked subpartition, since the sets S_w are disjoint and of size at least b , and

$$|L| = (n - 1)/((\varepsilon n/k) - 1) \geq (n/(\varepsilon n/k)) = k/\varepsilon > k.$$

Otherwise, v^* is labeled with some downward-reachable $u \in L$. Our second case, which is also straightforward, is that after this labeling of the root, all sets S_w for $w \in L$ have size exactly b , then we have a (k, ε) -linked subpartition.

If not, then we are in the third case: v^* is labeled with some downward-reachable $u \in L$, and after this labeling there still exist downward-reachable nodes w that we have placed in L that do not have associated sets S_w of size b . We therefore need to prune our set L to a smaller set that has $|S_w| \geq b$ for each $w \in L$. The goal is to show that the smaller L we end up with still has enough elements; if that holds, then we have a (k, ε) -linked subpartition.

To show this, we proceed as follows. We say that $w \in L$ is *active* if $|S_w| < b$. We first observe that any active w must be downward-reachable from the root v^* . Indeed, if w is active and not downward-reachable from the root, then there is a $v \in L$ such that w is a descendent of v . But in the step when we placed v in L , it was not possible to label v with w , and hence we must have had $|S_w| = b$ at that point.

Next we claim that there are $< k$ active $w \in L$. To prove this, for each active w , we associate w with a leaf that is a descendent of w . (This can be w itself if w is a leaf.) Observe that the same leaf x cannot be associated with two distinct active w, w' , for then on the path from v^* to x , one of w or w' would be closer to v^* , and the other would not be downward-reachable from v^* . Given that we can associate a distinct leaf to each active w , and there are $< k$ leaves, there are $< k$ active $w \in L$.

We say that a node $w \in L$ is *inactive* if it is not active; that is, if $|S_w| = b$. Let L_0 be the inactive nodes of L and L_1 be the active nodes of L . We have

$$|L_0| + \sum_{w \in L_0} |S_w| + |L_1| + \sum_{w \in L_1} |S_w| = n.$$

We know that $|L_1| < k$ and $|S_w| < b$ for each $w \in L_1$; hence, using the fact that $\varepsilon = 1/2$, we have

$$|L_1| + \sum_{w \in L_1} |S_w| < k + kb = k(b + 1) = kn/(2k) = n/2.$$

It follows that $|L_0| + \sum_{w \in L_0} |S_w| > n/2$. But since $|S_w| = b$ for each $w \in L_0$, we have

$$n/2 < |L_0| + \sum_{w \in L_0} |S_w| = |L_0|(1 + b) = |L_0|n/(2k),$$

from which it follows that $|L_0| > k$. We now conclude the construction by declaring L to be L_0 ; since the sets S_w for $w \in L_0$ are all pairwise disjoint and each has size at least b , we have the desired (k, ε) -linked subpartition.

5. CONCLUSION AND FUTURE WORK

We have introduced a new line of inquiry for the envy-free and proportional cake cutting problems by considering local notions of fairness. We show interesting relations between these local fairness concepts and their global analogues. Besides introducing this new model, our main contribution has been to fully classify the class of graphs for which there is an oblivious single-cutter protocol for computing locally envy-free allocations. Furthermore, we quantify the degradation in welfare resulting from adding the local envy-freeness constraint on the allocations; in particular, we show that the known $\Omega(\sqrt{n})$ lower-bound for the (global) price of envy-freeness continues to hold even for sparse graphs.

It is of interest to give efficient protocols for computing locally envy-free allocations on rich classes of graphs without the single-cutter constraint, which we hope this work will inspire. Since local envy-freeness is a stronger condition than local proportionality, the same problem can also be considered for locally proportional allocations, where perhaps progress can be made for a broader class of graphs. Finally, whether there is a similar lower bound of $\Omega(\sqrt{n})$ for the price of local proportionality is an open question. Currently, the upper bound for both local fairness concepts is the loose $n - 1/2$ bound, and giving tighter bounds is another direction.

Acknowledgement

The first author would like to thank Felix Fischer for the introduction to the cake cutting problem. We would like to thank Rahmtin Rotabi for helpful insights for the proof of Theorem 18, Professors Barabra Grosz and Professor Radhika Nagpal whose courses on multiagent systems inspired the formulation of local fairness concepts, and Adam Brauer and Ofra Amir for helpful discussions during these courses. This work was supported in part by a Simons Investigator Award.

APPENDIX

A. TAKING ADVANTAGE OF IRREVOCABLE ADVANTAGE

Given a partial envy-free allocation, an agent i is said to *dominate* an agent j , if i remains envy-free of j even if the entire residue (the remaining subinterval of the cake) is allocated to j . We can thus define:

Definition 19. A domination graph on n is a graph where V is the set of agents and there is a directed edge (i, j) if i has irrevocable advantage over j .

Constraints on the number of agents each agent must dominate at a certain stage in the protocol can be used to extend partial envy-free allocations to complete ones. One salient example is the Aziz-Mackenzie protocol for K_4 , where they obtain a partial envy-free allocation using what they call the *Core Protocol* and the *Permutation Protocol* such that each agent is guaranteed to dominate at least two other agents. They then use the Post-Double Domination Protocol to extend this to a complete envy-free allocation. We generalize this protocol to any n . That is, given a partial envy-free allocation such that each agent dominates $n - 2$ other agents, we can apply Protocol 1 to extend the allocation to a complete allocation. This provides an alternate proof to a recent paper by Segal-Halevi et al. [26]. We first define a special class of graphs.

Definition 20. A pseudoforest is a graph where each vertex has at most one outgoing edge.

Each component of a pseudoforest is a subgraph of a cone of a DAG; since each node has at most one outgoing edge, there are at most n edges. If there are fewer than n , then it is a DAG. If there are exactly n , then there exists a cycle. Find this cycle and remove an edge $e = (u, v)$ from the cycle. The resulting graph will be a DAG, and we can therefore Protocol 1 by setting $u = c$. This makes the protocol key to extending partial envy-free allocations to complete ones under particular domination criteria.

Lemma 21. We can extend a partial globally envy-free allocation in which each agent dominates at least $n - 2$ other agents to a complete, envy-free allocation.

PROOF. Suppose we have a partial globally envy-free allocation (P_1, P_2, \dots, P_n) , with residue R . If each agent dominates at least $n - 2$ other agents, then the complement of the domination graph, denoted by G^c , is a pseudoforest. We apply Protocol 1 on G^c using residue R and denote this allocation by (R_1, R_2, \dots, R_n) . We want to show that $(P_1 \cup R_1, P_2 \cup R_2, \dots, P_n \cup R_n)$ is a globally envy-free allocation. Suppose it is not. Then, there exists i, j such that $V_i(P_i \cup R_i) < V_i(P_j \cup R_j)$, but this is only possible if either $V_i(P_i) < V_i(P_j)$ or $V_i(R_i) < V_i(R_j)$.

The assumption that each agent dominates at least $n - 2$ agents is necessary. In particular, suppose that there exists one agent that dominates only $n - 3$ other agents, such as in Example 22.




Figure 3: Counterexample to the extension lemma.

Example 22. Suppose that each agent $i \neq 2$ dominates every other agent but agents $i + 1 \pmod n$ and that agent 2 dominates agents $\{4, 5, \dots, n\}$. The domination graph is given in Figure 3. The complement of the domination graph is the cycle graph $(1, 2, 3, \dots, n, 1)$ plus the edge $(2, 1)$. It therefore consists of more than one simple cycle, and hence a direct application of Protocol 1 to the complement of the domination graph will not extend a partial allocation to a complete allocation.

REFERENCES

- [1] Noga Alon. Splitting necklaces. *Advances in Mathematics*, 63(3):247–53, 1987.
- [2] George Akerlof. Social distance and social decisions. *Econometrica*, 65(2):1005–1028, 1997.
- [3] Haris Aziz and Simon Mackenzie. A discrete and bounded envy-free cake cutting protocol for four agents. In *Proceedings of the Fourty-eighth Annual ACM Symposium on Theory of Computing*, STOC’16, pages 454–464, Cambridge, MA, USA, 2016. ACM.
- [4] Haris Aziz and Simon Mackenzie. A discrete and bounded envy-free cake cutting protocol for any agents. In *Proceedings of the 57th Annual Symposium on Foundations of Computer Science*, FOCS’2016, pages 416–427, New Brunswick, New Jersey, USA, 2016.
- [5] Simina Brânzei, Ariel Procaccia, and Jie Zhang. Externalities in cake cutting. *Proceedings of the 23rd International Joint Conference on Artificial Intelligence*, IJCAI’12, pages 55–61, Beijing, China, 2013.
- [6] Joseph Blatt, Stephanie Stewart, and Sol Garfunkel. More equal than others: Weighted voting ; zero sum games: Games of conflict. Technical report, Annenberg/CPB Project, 1986.
- [7] Steven Brams and Alan Taylor. An envy-free cake division protocol. *American Mathematical Monthly*, 102(1):9–18, 1995.
- [8] Steven Brams and Alan Taylor. *Fair Division: From Cake Cutting to Dispute Resolution*. Cambridge University Press, Cambridge, 1996.
- [9] Simina Brazei, Ariel Procaccia, and Jie Zhang. Externalities in cake cutting. *Proceedings of the 23rd International Joint Conference on Artificial Intelligence*, IJCAI’12, pages 55–61, Beijing, China, 2013.
- [10] Ronald Burt. The shadow of other people: socialization and social comparison in marketing. In S. Wuyt, M. Dekimpe, E. Gijbrecchts, and R. Peters, editors, *The Connected Customer*. Taylor and Francis, 2010.
- [11] Ioannis Caragiannis, Christis Kaklamanis, Panagiotis Kanellopoulos, and Maria Kyropoulou. The efficiency of fair division. *newblock Theory of Computing Systems*, 50, May 2012.
- [12] Yiling Chen, John Lai, David Parkes, and Ariel Procaccia. Truth, justice, and cake cutting. *Games and Economic Behavior*, 77, January 2013.
- [13] Yann Chevaleyre, Ulle Endriss, and Nicolas Maudet. Allocating goods on a graph to eliminate envy. In *Proceedings of AAAI-07*, pages 700–705, 2007.
- [14] Lester Dubins and E.H. Spanier. how to cut a cake fairly. *American Mathematical Monthly*, 68(1):1–17, 1961.
- [15] S. Evan and A. Paz. A note on cake cutting. *Discrete Applied Mathematics*, 7(3):285–96, 1984.
- [16] Leon Festinger. A theory of social comparison process. *Human Relations*, 7(2), 1954.
- [17] Michael Kearns. Graphical games. In Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay Vazirani, editors, *Algorithmic Game Theory*, chapter 7. Cambridge University Press, 2007.
- [18] B. Knaster. Sur le probleme du partage pragmatique de h. steinhaus. *Annales de la Societe Polonaise de Mathematique*, 19:228–230, 1946.
- [19] David Kurokawa, John Lai, and Ariel Procaccia How to cut a cake before the party ends. In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence*, AAAI’13, Bellevue, WA, USA, 2013.
- [20] Malik Magdon-Ismaïl, Costas Busch, and Mukkai Krishnamoorthy. Cake-cutting is not a piece of cake. In *Proc. 20th Annual Symposium on Theoretical Aspects of Computer Science*, pages 596–607, 2003.
- [21] Michael McBride. Relative-income effects on subjective well-being in the cross section. *Journal of Economic Behavior and Organization*, 45(3): 251–278, 2001.
- [22] Ariel Procaccia. Thou shalt covet thy neighbor’s cake. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence*, IJCAI’09, pages 239–244, San Francisco, CA, USA, 2009. Morgan Kaufmann Publishers Inc.
- [23] Ariel Procaccia. Cake cutting. *Communications of the ACM*, 56(7), 2013.
- [24] Ariel Procaccia. Cake cutting algorithms. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A.D. Procaccia, editors, *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
- [25] Jack Robertson and William Webb. *Cake Cutting Algorithms: Be Fair If You Can*. A.K. Peters Ltd., London, 1998.
- [26] Erel Segal-Halevi, Avinatan Hassidim, and Yonatan Aumann. Waste makes haste: bounded time algorithms for envy-free cake cutting with free disposal. In *ACM Transactions of Economics and Computation*, 13(1), 2016.
- [27] Hugo Steinhaus. The problem of fair division. *Econometrica*, 16:101–104, 1948.
- [28] Walter Stromquist. Envy-free cake divisions cannot be found by finite protocols. *The Electronic Journal of Combinatorics*, 15(R11), 2008.