Constraints in Fair Division

WARUT SUKSOMPONG
National University of Singapore

The fair allocation of resources to interested agents is a fundamental problem in society. While the majority of the fair division literature assumes that all allocations are feasible, in practice there are often constraints on the allocation that can be chosen. In this survey, we discuss fairness guarantees for both divisible (cake cutting) and indivisible resources under several common types of constraints, including connectivity, cardinality, matroid, geometric, separation, budget, and conflict constraints. We also outline a number of open questions and directions.

Categories and Subject Descriptors: J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics*

General Terms: Economics, Algorithms, Theory

Additional Key Words and Phrases: Constraints, Fair division, Social choice

1. INTRODUCTION

A fundamental problem in society is how to best allocate resources to interested agents given the agents' preferences over the resources. Indeed, the allocation of course slots to students, machine processing time to users, and personnel to organizations all fall under this general problem. The study of how to perform the allocation in a fair manner—often referred to as *fair division*—has received ongoing interest from researchers in mathematics, economics, and computer science alike [Brams and Taylor 1996; Moulin 2003; Thomson 2016; Walsh 2020].

The majority of work on fair division assumes that any allocation of the resource to the agents is feasible. In many applications, however, there are constraints on the allocation that can be made. For instance, when allocating offices in a university building to research groups, it is helpful for each group to receive a *connected* set of offices in order to facilitate communication within the group. Another example is the division of land among inhabitants: Since a thin or highly zigzagged piece of land is unlikely to have much value even if its area is large, the *geometric* shape of the allocated land must be taken into account. Additionally, if the land is used to grow crops, keeping different plots *separated* can help prevent cross-fertilization. Approaches that ignore such relevant constraints may lead to highly undesirable or even completely unusable solutions.

In this survey, we discuss several types of constraints that have been considered in the fair division literature, most of them in the past few years. Besides connectivity, geometric, and separation constraints, these include cardinality, budget, as well as conflict considerations. As will be evident throughout the survey, handling such constraints usually requires developing new fairness concepts and algorithmic approaches. In addition to outlining the most significant results in our view, we also highlight a number of open questions and directions for future work.

Author's address: warut@comp.nus.edu.sg

2. PRELIMINARIES

In the setting of fair division, there is a set of agents $N = \{1, 2, \ldots, n\}$ among whom we wish to divide a resource. The literature distinguishes between two main types of resource: *divisible*—such as land and time—and *indivisible*—such as books, paintings, and jewelry.

When the resource is divisible, the problem is commonly known as *cake cutting*, with the cake serving as a metaphor for the resource [Brams and Taylor 1996; Robertson and Webb 1998; Procaccia 2016]. The cake is represented by the interval $M = [0, 1]$, and each agent i has a utility function u_i over the cake. We denote agent i's utility for the subset of cake S by $u_i(S)$; when S is a single interval $[x, y]$, we simplify notation and write $u_i(x, y)$ instead of $u_i([x, y])$. The utility functions are nonnegative, nonatomic (i.e., the utility for any single point is 0), and normalized so that $u_i(0, 1) = 1$ for all $i \in N$. On the other hand, in the indivisible resource setting [Bouveret et al. 2016; Markakis 2017], there is a set of goods $M = \{g_1, \ldots, g_m\}$, and each agent i has a nonnegative utility $u_i(M')$ for each subset of goods $M' \subseteq M$. In both settings, the utility functions are typically assumed to be *monotonic*—adding extra cake or goods to a bundle cannot decrease an agent's utility for the bundle—and are moreover often assumed to be *additive*.

For cake cutting, an important issue is how an algorithm can access the agents' utility functions, since the functions are not discrete. This is typically done using a model of Robertson and Webb [1998], which allows two types of queries:

—$\textsc{Eval}_i(x, y)$: Output the utility $u_i(x, y)$.

—$\textsc{Cut}_i(x, \alpha)$: Output the leftmost point y such that $u_i(x, y) = \alpha$, or return that no such y exists.

Clearly, a polynomial-time algorithm in this model can only make a polynomial number of queries, but the converse does not necessarily hold. For computation in the indivisible goods setting, if the utilities are not additive, we assume a *utility oracle model*: an algorithm can query the value $u_i(M')$ for any $i \in N$ and $M' \subseteq M$.

An allocation is denoted by $A = (A_1, \ldots, A_n)$, where A_i is the bundle allocated to agent i and $A_i \cap A_j = \emptyset$ for all $i \neq j$; the allocation is said to be *complete* if it allocates the entire resource, that is, $A_1 \cup \cdots \cup A_n = M$. Unless specified otherwise, we assume that allocations must be complete.

In order to reason about fairness, we must define what it means for an allocation to be "fair". Three of the most prominent definitions of fairness are:

—**Envy-freeness:** An allocation A is said to be *envy-free* if $u_i(A_i) \geq u_i(A_j)$ for every pair of agents i and j.

—**Proportionality:** An allocation A is said to be *proportional* if $u_i(A_i) \geq \frac{1}{n} \cdot u_i(M)$ for every agent i.

—**Maximin share fairness:** The *maximin share* of agent i, denoted by MMS_i, is the maximum utility that i can obtain by partitioning the resource into n parts and receiving the worst part according to u_i. A partition such that every part yields utility at least MMS_i according to u_i is called agent i's *maximin partition*, and an allocation such that every agent i receives utility at least MMS_i is called a *maximin allocation*.

Note that the first notion is "envy-based" whereas the latter two are "share-based". Crucially, when there are constraints on the feasible allocation, the same constraints are imposed on the partitions in the calculation of agents' maximin shares. As we will see, the robustness encoded in the definition of maximin share fairness makes it a particularly suitable benchmark in several constrained settings.

For indivisible goods, a natural relaxation of envy-freeness is the following:

—**Envy-freeness up to k goods:** An allocation A of indivisible goods is said to be *envy-free up to k goods (EFk)* if for each pair of agents i and j, there exists a set $B \subseteq A_j$ with $|B| \leq k$ such that $u_i(A_i) \geq u_i(A_j \setminus B)$.

3. CONNECTIVITY CONSTRAINTS

The most frequently studied constraint in fair division is the *connectivity* of the resource allocated to each agent. Connectivity has been considered in cake cutting for several decades—as Stromquist [1980] memorably put it, without this constraint, there is a danger that agents who only hope for a modest piece of cake will instead receive a "countable union of crumbs". We survey work on connectivity in cake cutting as well as in extensions where the resource is represented by a graph.

3.1 Connected Cake Cutting

An early result in cake cutting is a protocol by Dubins and Spanier [1961] that computes a connected proportional allocation for any number of agents with additive utilities. A discrete version of their protocol proceeds as follows: First, ask each agent i for $\text{Cut}_i(0, 1/n)$, i.e., the leftmost point y_i such that $u_i(0, y_i) = 1/n$. An agent i^* who gives the smallest answer y_{i^*} receives the piece $[0, y_{i^*}]$; the process is then recursively repeated with the remaining $n - 1$ agents and remaining cake. When there is one agent left, that agent receives the entire leftover cake. The Dubins–Spanier protocol makes $O(n^2)$ queries in the Robertson–Webb model. Even and Paz [1984] reduced the number of queries to $O(n \log n)$ using a divide-and-conquer approach, and Edmonds and Pruhs [2011] established a matching lower bound of $\Omega(n \log n)$, even when connectivity is not required.

While a connected proportional allocation of the cake is easy to obtain, combining connectivity and envy-freeness turns out to be much more complicated.[1] Stromquist [1980] and Su [1999] showed via topological approaches that a connected envy-free allocation always exists. Despite this existence result, Stromquist [2008] remarkably proved that no algorithm can compute such an allocation using a finite number of queries when there are at least three agents.[2] This raises the question of whether we can efficiently compute a connected allocation with low envy. If we are only concerned with the number of queries made, the answer is positive. For any $\delta > 0$, we say that an allocation is *δ-additive-envy-free* if $u_i(A_i) \geq u_i(A_j) - \delta$ for all $i, j \in N$, and *δ-multiplicative-envy-free* if $u_i(A_i) \geq u_i(A_j)/\delta$ for all $i, j \in N$.

THEOREM 3.1 [BRÂNZEI AND NISAN 2017]. *For cake cutting with additive util-*

---

[1]Even without connectivity, currently the best envy-free protocol requires $n^{n^{n^{n^{n^n}}}}$ queries [Aziz and Mackenzie 2016].

[2]When $n = 2$, envy-freeness and proportionality are equivalent.

ities and any $\varepsilon > 0$, there exists an algorithm that computes a connected ε-additive-envy-free allocation using $O(n/\varepsilon)$ queries in the Robertson–Webb model.

Brânzei and Nisan's algorithm works by asking each agent to cut the cake into pieces of value at most $\varepsilon/4$ and performing a brute-force search over the space of all contiguous allocations with respect to the union of all agents' cuts; the existence of an ε-additive-envy-free allocation in the search space is guaranteed by the aforementioned result of Stromquist [1980]. Despite its low query complexity, the algorithm runs in time exponential in n, even for constant ε. When polynomial time is required, the current best additive approximation is due to Goldberg et al. [2020].

THEOREM 3.2 [GOLDBERG ET AL. 2020]. *For cake cutting with additive utilities, there is a polynomial-time algorithm that computes a connected $1/3$-additive-envy-free allocation using $O(n^2)$ queries in the Robertson–Webb model.*

Goldberg et al.'s algorithm shares similar ideas with the Dubins–Spanier protocol. First, using the $\text{CUT}_i(0, 1/3)$ query, it asks each agent i for the leftmost point y_i such that $u_i(0, y_i) = 1/3$. If such a point exists for at least one agent, an agent i^* who gives the smallest answer y_{i^*} receives the piece $[0, y_{i^*}]$, and the algorithm recurses over the remaining $n - 1$ agents and remaining cake. Otherwise, if such a point y_i exists for none of the agents, the remaining cake is allocated to an arbitrary remaining agent. If we have run out of agents but there is still cake left, the leftover cake is allocated to the agent who receives the last piece. To see why the algorithm yields the desired guarantee, observe that an agent who receives a piece of value $1/3$ has utility $2/3$ for the remaining cake, and therefore has additive envy at most $1/3$. On the other hand, the algorithm ensures that any agent who receives utility less than $1/3$ does not value any other agent's piece more than $1/3$, so the additive envy of such an agent is also at most $1/3$.

Although Goldberg et al.'s algorithm achieves a decent additive approximation of envy-freeness, it may leave certain agents empty-handed, leading to an unbounded multiplicative approximation. Arunachaleswaran et al. [2019] proposed a different algorithm that simultaneously attains both multiplicative and additive guarantees.

THEOREM 3.3 [ARUNACHALESWARAN ET AL. 2019]. *For cake cutting with additive utilities and any constant $\delta \in (0, 1/3]$, there exists a polynomial-time algorithm that computes a connected $(2 + \frac{9\delta}{n})$-multiplicative-envy-free allocation.*[3]

Note that a 2-multiplicative-envy-free allocation is also 1/3-additive-envy-free; this implies that Arunachaleswaran et al.'s algorithm guarantees an additive envy of close to 1/3. Their algorithm is more complex than Goldberg et al.'s and requires allocating successively higher-valued pieces to agents while ensuring that the partial allocation at any point during the execution is approximately envy-free.

In light of these algorithms, an important question is whether the approximation guarantees can be improved while at the same time maintaining efficient computation.

[3]The conference version of their paper achieved a weaker bound of $3 + \frac{9\delta}{n}$; the stronger bound was provided in their arXiv version.

Open problem 3.4. *For cake cutting with additive utilities, which values of ε admit a polynomial-time algorithm that computes a connected ε-additive- or ε-multiplicative-envy-free allocation? Does there exist a polynomial-time approximation scheme (PTAS) or a fully polynomial-time approximation scheme (FPTAS) for either type of approximation?*[4]

Certain applications require additional constraints to be imposed on top of connectivity. For instance, there could be a temporal ordering in which the agents must be served, perhaps due to notions of seniority or the ease of switching from one agent to another in the service. Another example is when we divide a parcel of land and there is a road crossing the parcel, so one of the cut points must fall at the position of the road. A connected envy-free allocation is no longer guaranteed to exist when extra constraints are imposed. Goldberg et al. [2020] showed that for several constraints, deciding whether such an allocation exists is NP-hard.

3.2 Indivisible Goods on a Graph

In contrast to cake cutting, the items in indivisible resource allocation do not come with an inherent linear order. Nevertheless, there are scenarios in which such an order is present, for example when allocating offices along a corridor or assigning hourly time slots for using a facility. Bouveret et al. [2017] proposed a general model wherein the items correspond to the vertices on an undirected graph G—this allows us to capture complex spatial or temporal relationships among the items. These authors showed that when the graph is a tree, maximin share fairness is an appropriate fairness notion. Recall that connectivity is also imposed when calculating the agents' maximin shares.

Theorem 3.5 [Bouveret et al. 2017]. *For indivisible goods allocation with monotonic utilities, if G is a tree, there exists a connected maximin allocation.*[5] *Moreover, if the utilities are additive, such an allocation can be found in polynomial time, and each agent's maximin share can also be computed in polynomial time.*

To understand how Bouveret et al.'s allocation algorithm works, consider first the case where G is a path. The algorithm proceeds over the path from left to right in a similar manner as the Dubins–Spanier protocol: At each step, for each agent i, it identifies a good g^i such that the block of remaining goods up to g^i is worth at least MMS_i to the agent. The algorithm then chooses an agent i^* with the leftmost g^{i^*}, allocates the block up to g^{i^*} to her, and recurses on the remaining agents and goods. The correctness follows from the observation that every time a block of goods is allocated, it interferes with at most one additional part of each remaining agent's maximin partition, so at the end the last agent i is still left with an entire part of her partition, and with it utility at least MMS_i.

When G is a tree, there is no notion of "left" or "right", so it is not immediate how to generalize the aforementioned algorithm. Nevertheless, a slight reformulation of

[4] Deng et al. [2012] gave an FPTAS for additive approximation when $n = 3$.

[5] Bouveret et al. [2017] stated their result for additive utilities, but their proof also works for general monotonic utilities. Furthermore, the proof can be extended to (not necessarily connected) acyclic graphs.

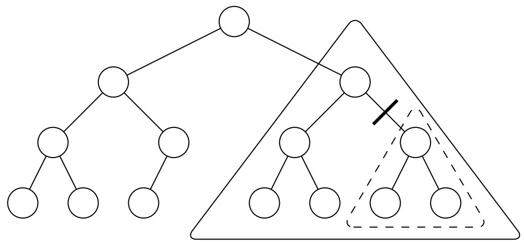

Fig. 1. Illustration of the last diminisher procedure on a tree. If the solid triangle is the current bundle and an agent diminishes it according to the thick mark, the new bundle is the dashed triangle.

the Dubins–Spanier protocol that turns out to be helpful is as a *last diminisher* procedure: every agent has an opportunity to "diminish" the existing piece whenever a smaller piece would still yield utility $1/n$ to her, and the last agent to diminish receives the corresponding piece. Given indivisible goods on a tree, Bouveret et al.'s algorithm starts by rooting the tree at an arbitrary vertex. Each agent i is allowed to diminish the tree if a subtree of it already yields utility at least MMS_i (see Figure 1). The last agent to diminish receives the subtree, and the algorithm recurses over the remaining agents and leftover tree. The proof of correctness is similar to the one for paths in the previous paragraph.

What happens if G is not a tree? Bouveret et al. [2017] also gave an example demonstrating that even when G is a simple cycle, the maximin share guarantee can no longer be made. Consider 8 goods and 4 agents with the following utilities:

	g_1	g_2	g_3	g_4	g_5	g_6	g_7	g_8
Agents 1, 2	1	4	4	1	3	2	2	3
Agents 3, 4	4	4	1	3	2	2	3	1

The goods $g_1, \dots, g_8$ lie on a cycle in this order. All agents have a maximin share of 5: For the first two agents, this can be seen by the partition

$$\{g_1, g_2\}, \{g_3, g_4\}, \{g_5, g_6\}, \{g_7, g_8\},$$

while for the last two agents, this is witnessed by the partition

$$\{g_2, g_3\}, \{g_4, g_5\}, \{g_6, g_7\}, \{g_8, g_1\}.$$

However, one can check that no connected allocation gives every agent a utility of at least 5; in fact, since all utilities are integers, one of the agents must receive utility at most 4. Lonc and Truszczynski [2020] proved that when agents have additive utilities and G is a cycle, it is possible to guarantee to every agent $\frac{\sqrt{5}-1}{2} \approx 0.62$ times her maximin share, and better bounds can be obtained when the number of agents is small.

Open problem 3.6. *For indivisible goods allocation, for each graph G and number of agents n, what is the maximum r such that for any n agents with additive/monotonic utilities, there exists a connected allocation that gives each agent i a utility of at least $r \cdot \text{MMS}_i$?*

Notice that although a cycle permits more connected allocations than a path, it offers less guarantee in terms of the maximin share—this is a consequence of the connectivity constraint being imposed on the maximin share calculation. Bei et al. [2021] studied the *price of connectivity*, which they defined as the worst-case ratio between the unconstrained and constrained maximin shares for each G and n.

Unlike in cake cutting, it is clear that envy-freeness and proportionality cannot always be attained with indivisible goods—the canonical example involves two agents who compete for a single valuable good. Bouveret et al. [2017] proved that deciding whether an envy-free allocation exists is NP-hard when G is a path or a star, whereas the corresponding question with proportionality remains hard for paths but can be solved in polynomial time for stars. Bouveret et al. [2019] performed similar complexity analyses when the items are *chores* (i.e., they yield disutilities to the agents).

In order to bound the envy among agents, Bilò et al. [2019] focused on envy-freeness relaxations.

THEOREM 3.7 [BILÒ ET AL. 2019]. *For indivisible goods allocation with monotonic utilities, if G is a path, there exists a connected EF2 allocation.*[6]

Bilò et al.'s proof relies on a variant of Sperner's lemma; a similar technique was used by Su [1999] to establish the existence of a connected envy-free allocation in cake cutting. Bilò et al. also showed that EF2 can be improved to EF1 when $n \leq 4$. Suksompong [2019] demonstrated that the same improvement can be attained when agents have *binary* additive utilities, i.e., each agent has utility 0 or 1 for each good. For constant n, an allocation that meets any of these guarantees can be found efficiently by brute-force search, since the number of connected allocations is polynomial in m in this case.

OPEN PROBLEM 3.8. *For indivisible goods allocation with additive/monotonic utilities, if G is a path, does there always exist a connected EF1 allocation? What is the complexity of computing a connected EF2 allocation? What about computing a connected EF1 allocation when agents have binary additive utilities?*

Misra et al. [2021] considered another fairness notion called *equitability* with respect to paths. Beyond paths, Bilò et al. [2019] characterized the graphs that guarantee EF1 when $n = 2$. Bei et al. [2021] generalized this result by providing for each graph G the smallest k such that a connected EFk allocation with respect to G always exists. Igarashi and Peters [2019] studied fairness in conjunction with the economic efficiency notion of Pareto optimality.

OPEN PROBLEM 3.9. *For indivisible goods allocation, for each graph G and number of agents n, what is the minimum k such that for any n agents with additive/monotonic utilities, there exists a connected EFk allocation?*

3.3 Graphical Cake Cutting

Like indivisible goods, the interval cake in cake cutting can also be generalized to an arbitrary graph: the resulting model, called *graphical cake cutting*, captures the

[6]The theorem extends to all graphs G that contain a Hamiltonian path, since finding a connected EF2 allocation with respect to G reduces to finding one with respect to its Hamiltonian path.

division of graphical resources such as road networks. While full proportionality can no longer be guaranteed together with connectivity, Bei and Suksompong [2021] showed that more than half of the guarantee can be recovered.

THEOREM 3.10 [BEI AND SUKSOMPONG 2021]. *For graphical cake cutting with additive utilities and any connected graph G, there exists a connected allocation such that every agent receives utility at least*[7] $\frac{1}{2n-1}$.

The tightness of this bound follows by considering a star with $2n-1$ edges and uniform utilities over the edges: in any connected allocation, one of the agents is restricted to at most one edge. Bei and Suksompong's algorithm uses a recursive approach similar to the Dubins–Spanier protocol. Letting $\alpha = \frac{1}{2n-1}$, at each stage, it finds a connected piece worth at least α to one agent and at most 2α to every remaining agent, such that the cake with this piece taken away remains connected; this ensures that at the end, the last agent receives utility at least $1-(n-1)\cdot 2\alpha = \alpha$. In order to identify a desired piece, the algorithm roots the tree at an arbitrary vertex and identifies a lowest vertex v whose subtree is worth at least α to some agent. If at least one of the branches of v along with the corresponding subtree is worth at least α to some agent, the algorithm finds the lowest point on the branch such that the cake below it yields utility exactly α to some agent—this piece of cake is therefore worth at most α to every other agent. On the other hand, if none of the branches along with its subtree is worth at least α to any agent, the algorithm adds these branches one by one until the resulting cake yields utility at least α to some agent—this guarantees that the cake has value at most 2α to every other agent.

For two agents, Bei and Suksompong determined the optimal guarantee that can be made for each graph; perhaps surprisingly, this guarantee is always either $1/2$ or $1/3$. Nevertheless, little is known beyond the case $n = 2$.

OPEN PROBLEM 3.11. *For graphical cake cutting with additive utilities, for each graph G and number of agents $n \geq 3$, what is the largest c such that there always exists a connected allocation that gives each agent a utility of at least c?*

Graphical cake cutting has also been studied with respect to other fairness notions. Elkind et al. [2021a] showed that maximin share fairness can be attained for acyclic graphs if sharing vertices is prohibited, but may be unattainable even for stars if sharing is allowed; this further highlights the difference that the sharing assumption makes. Igarashi and Zwicker [2021] considered envy-freeness when vertices cannot be shared and characterized the graphs that guarantee the existence of a connected envy-free allocation simultaneously for all n.

Given that a motivation for imposing connectivity is to avoid allocating a "union of crumbs", a natural question is what happens when we relax this constraint and allow each agent to receive a small number of connected pieces. Bei and Suksompong [2021] showed that when $n = 2$, if the two agents can receive up to a total of $d + 1$ connected pieces, a utility of $\frac{1}{2} - \frac{1}{2\cdot 3^d}$ can be guaranteed for each

[7] A subtle point is that for this result, we must assume that vertices can be shared by more than one agent. Otherwise, for a star with $t \gg n$ edges and uniform utilities over the edges, each of the $n-1$ agents who do not receive the center vertex will be restricted to at most one edge and therefore receive utility at most $1/t \ll 1/n$.

agent, and this bound is tight. Exploring this kind of relaxations further for both cake cutting and indivisible goods is an interesting direction for future work.[8]

4. CARDINALITY AND MATROID CONSTRAINTS

Imagine a museum that is faced with the task of distributing a set of exhibits among its newly opened branches. Since the branches have a capacity limit, it wants to ensure that the allocation is *balanced*, i.e., the numbers of exhibits that different branches receive differ by at most one. Can the museum always find a balanced allocation that is fair to the curators of these new branches?

If the curators' utilities over the exhibits are additive, the task can be accomplished by the classic *round-robin algorithm*, wherein the branches take turns picking their favorite exhibit from the remaining ones until the exhibits run out. The output of the round-robin algorithm is always EF1. However, this guarantee relies crucially on additivity.[9] For arbitrary monotonic utilities, Kyropoulou et al. [2020] showed that a balanced EF1 allocation exists when there are two agents, but the question remains intriguingly open for higher numbers of agents.

Open problem 4.1. *For indivisible goods allocation with monotonic utilities and $n \geq 3$, does there always exist a balanced EF1 allocation?*

Balancedness is a basic cardinality constraint.[10] More generally, the exhibits may be categorized into paintings, sculpture, pottery, and so on, and the museum wants to ensure a balanced distribution within each category. When utilities are additive, these sophisticated constraints can still be satisfied.

Theorem 4.2 [Biswas and Barman 2018]. *For indivisible goods allocation, if the agents have additive utilities and the goods are categorized into types, there exists an EF1 allocation that is balanced with respect to each type. Moreover, such an allocation can be found in polynomial time.*

Biswas and Barman's algorithm works by allocating the goods one category at a time using round-robin. After it has allocated all goods of a category, it updates the "envy graph" between agents by eliminating cycles in the graph.[11] The round-robin order for the next category of goods is then determined by an arbitrary topological ordering in the resulting acyclic envy graph—in particular, if agent i envies agent j, then i comes before j in the order.

Besides EF1, Biswas and Barman also showed that a constant fraction of the agents' maximin shares can be guaranteed under category constraints; their approximation was improved by Hummel and Hetland [2021b]. Moreover, Biswas and Barman considered a broader class of constraints defined by a matroid. Specifically, the bundle of goods allocated to each agent must form an independent set of the

[8]Similar relaxations have also been studied in cake cutting by Arunachaleswaran and Gopalakrishnan [2018] and Segal-Halevi [2021].
[9]In fact, the guarantee holds for the larger class of *responsive utilities*; see, e.g., [Kyropoulou et al. 2020].
[10]Another application in which balancedness plays an important role is assigning conference papers to reviewers—see the paper by Garg et al. [2010] for an example of work in this line.
[11]This idea was used by Lipton et al. [2004] in the unconstrained allocation setting.

(common) matroid. They showed that when the matroid is laminar and the agents have identical utilities, EF1 can again be guaranteed.

THEOREM 4.3 [BISWAS AND BARMAN 2018]. *For indivisible goods allocation and a laminar matroid $\mathcal{M}$, if the agents have identical additive utilities, each agent's bundle must form an independent set of $\mathcal{M}$, and there is at least one feasible allocation, then there is a feasible EF1 allocation. Moreover, such an allocation can be found in polynomial time.*[12]

The conference version of Biswas and Barman's paper made this claim for all matroids. However, their proof contained a flaw, and the statement was subsequently weakened to laminar matroids in their arXiv version. The existence question therefore remains open, even for two agents with identical utilities.

OPEN PROBLEM 4.4. *For indivisible goods allocation with identical additive utilities and a matroid $\mathcal{M}$, if each agent's bundle must form an independent set of $\mathcal{M}$ and there is at least one feasible allocation, does there always exist a feasible EF1 allocation? What about for non-identical additive utilities?*

Dror et al. [2021] examined EF1 in the setting where different agents may have different matroid constraints. Gourvès et al. [2014] and Gourvès and Monnot [2019] also considered matroid constraints, but imposed the constraint on the set of all allocated goods (as opposed to each agent's set); their setting therefore allows allocations to be incomplete. Li and Vetta [2021] studied maximin share approximations for set systems satisfying the *hereditary property*, i.e., any subset of a feasible set is also feasible—note that all matroids satisfy this property.

In cake cutting, an analog of cardinality constraints where all agents must receive the same amount of cake has been studied by Jojić et al. [2021].

5. GEOMETRIC CONSTRAINTS

When allocating two-dimensional resources such as land or advertising spaces, geometric considerations are crucial: it is hard to build a house on a 5×500 meter land plot or put an advertisement on a 1 centimeter wide strip. Segal-Halevi et al. [2017] studied land division with the assumption that each piece must be a rectangle whose length-to-width ratio is at most r, for some given parameter $r \geq 1$. When $r = 1$, the pieces are required to be squares. Similarly to graphical cake cutting (Section 3.3), full proportionality cannot be attained under such constraints, but a constant factor can be recovered.

THEOREM 5.1 [SEGAL-HALEVI ET AL. 2017]. *For land division with additive utilities, if the land is a square, there exists an allocation such that all allocated pieces are squares and every agent receives utility at least $\frac{1}{4n-4}$.*

The algorithm that achieves this guarantee is formulated as a sequence of auctions, in which agents bid for pieces that they view as having sufficiently high value. On the other hand, Segal-Halevi et al. proved an upper bound of $\frac{1}{2n}$, which is tight for $n = 2$ but leaves a gap for larger n.

[12]This assumes, as is standard in matroid theory, that an algorithm can query whether a set of goods is independent in constant time.

Open problem 5.2. *For land division with additive utilities and each number of agents n, if the land is a square, what is the largest c for which there always exists an allocation such that all allocated pieces are squares and every agent receives a utility of at least c? What if instead of being squares, the pieces are required to have a length-to-width ratio of at most r for some given $r \geq 1$?*

Land allocation with pieces of usable shapes has also been studied with respect to envy-freeness [Segal-Halevi et al. 2020] and maximin share fairness [Elkind et al. 2021b].

6. SEPARATION CONSTRAINTS

In times of a pandemic, a ubiquitous restriction concerns social distancing: each pair of agents should be separated by at least a certain distance s, where $s > 0$ is a given parameter. Separation constraints are also relevant in the allocation of machine processing time, where we need time to erase data from the previous process before the next one can be started, as well as in land division, where we want space between different plots in order to avoid cross-fertilization. In cake cutting, envy-freeness and proportionality cannot always be satisfied in light of separation, for example when all agents place their entire value on a tiny piece of length less than s. Moreover, unlike with graphical or geometric constraints (Sections 3.3 and 5, respectively), the same example shows that even approximate proportionality cannot necessarily be fulfilled. Fortunately, maximin share fairness once again comes to the rescue.

Theorem 6.1 [Elkind et al. 2021c]. *For cake cutting with monotonic utilities and separation constraints, there exists a connected maximin allocation. Moreover, if the utilities are additive, given the maximin shares of all agents, there exists a polynomial-time algorithm that computes such an allocation using $O(n^2)$ queries in the Robertson–Webb model.*

Elkind et al.'s algorithm processes the cake from left to right in a similar way as the Dubins–Spanier protocol. An interesting question is whether it is possible to use only $O(n \log n)$ queries, as the Even–Paz protocol does in the setting without separation (cf. Section 3.1).

Open problem 6.2. *For cake cutting with additive utilities and separation constraints, given the maximin shares of all agents, is there an algorithm that computes a connected maximin allocation using $O(n \log n)$ queries in the Robertson–Webb model?*

The algorithm in Theorem 6.1 requires the knowledge of agents' maximin shares in order to make the appropriate cut queries. Elkind et al. proved that, perhaps surprisingly, this knowledge cannot be attained with a finite number of queries in the Robertson–Webb model; the impossibility holds even when $n = 2$. They therefore considered approximation and showed that for any $\varepsilon > 0$, a connected allocation in which each agent i receives utility at least $\text{MMS}_i - \varepsilon$ can be found using $O(n^2 \log(1/\varepsilon))$ queries.

In addition to cake cutting, Elkind et al. [2021c] investigated *pie cutting*, where the pie represents a one-dimensional circular resource such as the shoreline of an island or a daily schedule for using a facility. For pie cutting, maximin share fairness or any multiplicative approximation thereof cannot be attained. Nevertheless, a relaxation that turns out to be appropriate is the 1-*out-of-*k *maximin share*, where we compute an agent's maximin share by partitioning into k parts instead of n, for a given parameter $k > n$—for pie cutting, it suffices to take $k = n + 1$. The same authors also examined graphical cake cutting [Elkind et al. 2021a] and land division [Elkind et al. 2021b] under separation constraints; in both of these settings, the 1-out-of-k maximin share provides useful fairness guarantees as well.

7. BUDGET CONSTRAINTS

The cardinality constraints in Section 4 implicitly assume that all goods take up the same amount of space. A natural generalization studied by Wu et al. [2021] is to allow different goods to have varying *costs*, which can represent the amount of space that they take, and endow each agent with a limited *budget* for the goods. In the case that no complete allocation is budget-feasible, Wu et al. assumed that the remaining goods go to a *charity*. Since agents may have distinct budgets, the definition of EF1 must be adjusted accordingly: An allocation is said to be EF1 if for any pair of agents i, j, for any subset $X_j \subseteq A_j$ of cost at most i's budget, there exists a set $B \subseteq X_j$ with $|B| \leq 1$ such that $u_i(A_i) \geq u_i(X_j \setminus B)$. More generally, for any $\tau \in [0, 1]$, we can define τ-EF1 in the same way except that the inequality is changed to $u_i(A_i) \geq \tau \cdot u_i(X_j \setminus B)$. Wu et al. focused on the *maximum Nash welfare (MNW)* solution, which chooses an allocation that maximizes the product of the agents' utilities subject to budget-feasibility.

THEOREM 7.1 [WU ET AL. 2021]. *For indivisible goods allocation with additive utilities and budget constraints, the MNW solution is always* $1/4$*-EF1. Moreover, the factor* $1/4$ *is tight.*

Theorem 7.1 stands in contrast with the unconstrained setting, where Caragiannis et al. [2019] showed that the MNW solution is fully EF1. It leaves the question of whether other methods can achieve EF1 with budget constraints.

OPEN PROBLEM 7.2. *For indivisible goods allocation with additive utilities and budget constraints, does there always exist an EF1 allocation?*

When the agents have identical additive utilities, Gan et al. [2021] improved the approximation factor to $1/2$ via a greedy algorithm which runs in polynomial time.

8. CONFLICT CONSTRAINTS

If the goods to be allocated represent activities, it may not be possible for an agent to participate in activities whose time periods overlap. Hummel and Hetland [2021a] modeled such constraints using a *conflict graph*, an undirected graph in which there is an edge between two goods if and only if they cannot both be allocated to the same agent. Under mild conditions, these authors showed that a constant fraction of the maximin share can be guaranteed for any conflict graph.

THEOREM 8.1 [HUMMEL AND HETLAND 2021A]. *For indivisible goods allocation with additive utilities and a conflict graph, if the number of agents is strictly larger than the maximum degree of the conflict graph, then there exists an allocation respecting the conflict graph that gives each agent* i *a utility of at least* $\frac{1}{3} \cdot \mathrm{MMS}_i$.

Hummel and Hetland provided better bounds for certain values of n and conflict graphs, and also considered EF1. Chiarelli et al. [2020] studied conflict constraints for partial allocations that maximize the utility of the worst-off agent.

Another model involving conflict constraints was introduced by Hosseini et al. [2020]. In their model, there is a cake consisting of multiple layers, and an allocation is feasible if each agent receives a bundle such that the pieces belonging to different layers do not overlap. This model captures, for example, the assignment of time for using different facilities—an agent cannot use more than one facility simultaneously. Extending the existence of connected envy-free and proportional allocations of a single-layered cake to a multi-layered cake is nontrivial. Denote by k the number of layers of the cake.

THEOREM 8.2 [HOSSEINI ET AL. 2020]. *For multi-layered cake cutting with additive utilities, there exists a connected proportional allocation if k is a power of 2 and $n \geq k$. Moreover, there exists a connected envy-free allocation if $n = k = 2$.*

Igarashi and Meunier [2021] generalized the envy-freeness result to the case where n is a prime power and $n \geq k$. Like the existence results of Stromquist [1980] and Su [1999] for a single-layered cake, their result holds even for continuous utilities that are neither additive nor monotonic. Note that the condition $n \geq k$ is necessary, as there is no feasible way to allocate the entire cake otherwise. However, it remains open whether the existence guarantees hold for all $n \geq k$.

9. CONCLUSION

In this survey, we have outlined research on several types of constraints in fair division and highlighted a number of open questions. As we have seen, identifying appropriate notions and deriving guarantees for different types of constraints is a nontrivial task and often leads to intriguing technical questions. For instance, there is still a relatively large gap in polynomial-time envy-freeness approximation for connected cake cutting (Open problem 3.4), and whether an EF1 allocation of indivisible goods always exists under matroid constraints remains open even for two agents with identical utilities (Open problem 4.4).

In addition to obtaining stronger fairness guarantees under practical constraints, an important direction is to extend the study of constraints to more general settings; these include allocating a mix of indivisible and divisible goods [Bei et al. 2021] or a mix of goods and chores [Bogomolnaia et al. 2017; Segal-Halevi 2018; Aziz et al. 2019], handling agents with unequal entitlements [Farhadi et al. 2019; Cseh and Fleiner 2020; Babaioff et al. 2021; Chakraborty et al. 2021], and considering the allocation to groups instead of individual agents [Manurangsi and Suksompong 2017; Suksompong 2018; Segal-Halevi and Nitzan 2019]. Implementing the developed algorithms and making them publicly available, as was done for the unconstrained setting by the Spliddit website [Goldman and Procaccia 2014; Shah 2017], would help bring the theory of fair division closer to practice as well.

Acknowledgments

The author thanks Xiaohui Bei, Jiarui Gan, Ayumi Igarashi, Bo Li, Erel Segal-Halevi, and Inbal Talgam-Cohen for helpful comments, and acknowledges support from an NUS Start-up Grant.

REFERENCES

Arunachaleswaran, E. R., Barman, S., Kumar, R., and Rathi, N. 2019. Fair and efficient cake division with connected pieces. In *Proceedings of the 15th Conference on Web and Internet Economics (WINE)*. 57–70. Extended version available at CoRR abs/1907.11019.

Arunachaleswaran, E. R. and Gopalakrishnan, R. 2018. The price of indivisibility in cake cutting. *CoRR abs/1801.08341*.

Aziz, H., Caragiannis, I., Igarashi, A., and Walsh, T. 2019. Fair allocation of indivisible goods and chores. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*. 53–59.

Aziz, H. and Mackenzie, S. 2016. A discrete and bounded envy-free cake cutting protocol for any number of agents. In *Proceedings of the 57th Annual Symposium on Foundations of Computer Science (FOCS)*. 416–427.

Babaioff, M., Nisan, N., and Talgam-Cohen, I. 2021. Competitive equilibrium with indivisible goods and generic budgets. *Mathematics of Operations Research 46,* 1, 382–403.

Bei, X., Igarashi, A., Lu, X., and Suksompong, W. 2021. The price of connectivity in fair division. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*. 5151–5158.

Bei, X., Li, Z., Liu, J., Liu, S., and Lu, X. 2021. Fair division of mixed divisible and indivisible goods. *Artificial Intelligence 293*, 103436.

Bei, X. and Suksompong, W. 2021. Dividing a graphical cake. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*. 5159–5166.

Bilò, V., Caragiannis, I., Flammini, M., Igarashi, A., Monaco, G., Peters, D., Vinci, C., and Zwicker, W. S. 2019. Almost envy-free allocations with connected bundles. In *Proceedings of the 10th Innovations in Theoretical Computer Science Conference (ITCS)*. 14:1–14:21.

Biswas, A. and Barman, S. 2018. Fair division under cardinality constraints. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*. 91–97. Extended version available at CoRR abs/1804.09521.

Bogomolnaia, A., Moulin, H., Sandomirskiy, F., and Yanovskaya, E. 2017. Competitive division of a mixed manna. *Econometrica 85,* 6, 1847–1871.

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., Cechlárová, K., and Lesca, J. 2019. Chore division on a graph. *Autonomous Agents and Multi-Agent Systems 33,* 5, 540–563.

Bouveret, S., Chevaleyre, Y., and Maudet, N. 2016. Fair allocation of indivisible goods. In *Handbook of Computational Social Choice*, F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, Eds. Cambridge University Press, Chapter 12, 284–310.

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

Brânzei, S. and Nisan, N. 2017. The query complexity of cake cutting. *CoRR abs/1705.02946*.

Caragiannis, I., Kurokawa, D., Moulin, H., Procaccia, A. D., Shah, N., and Wang, J. 2019. The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation 7,* 3, 12:1–12:32.

Chakraborty, M., Igarashi, A., Suksompong, W., and Zick, Y. 2021. Weighted envy-freeness in indivisible item allocation. *ACM Transactions on Economics and Computation 9,* 3, 18:1–18:39.

Chiarelli, N., Krnc, M., Milanič, M., Pferschy, U., Pivač, N., and Schauer, J. 2020. Fair packing of independent sets. In *Proceedings of the 31st International Workshop on Combinatorial Algorithms (IWOCA)*. 154–165.

Cseh, A. and Fleiner, T. 2020. The complexity of cake cutting with unequal shares. *ACM Transactions on Algorithms 16,* 3, 29:1–29:21.

Deng, X., Qi, Q., and Saberi, A. 2012. Algorithmic solutions for envy-free cake cutting. *Operations Research 60,* 6, 1461–1476.

DROR, A., FELDMAN, M., AND SEGAL-HALEVI, E. 2021. On fair division under heterogeneous matroid constraints. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*. 5312–5320.

DUBINS, L. E. AND SPANIER, E. H. 1961. How to cut a cake fairly. *American Mathematical Monthly 68,* 1, 1–17.

EDMONDS, J. AND PRUHS, K. 2011. Cake cutting really is not a piece of cake. *ACM Transactions on Algorithms 7,* 4, 51:1–51:12.

ELKIND, E., SEGAL-HALEVI, E., AND SUKSOMPONG, W. 2021a. Graphical cake cutting via maximin share. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*. 161–167.

ELKIND, E., SEGAL-HALEVI, E., AND SUKSOMPONG, W. 2021b. Keep your distance: Land division with separation. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*. 168–174.

ELKIND, E., SEGAL-HALEVI, E., AND SUKSOMPONG, W. 2021c. Mind the gap: Cake cutting with separation. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*. 5330–5338.

EVEN, S. AND PAZ, A. 1984. A note on cake cutting. *Discrete Applied Mathematics 7,* 3, 285–296.

FARHADI, A., GHODSI, M., HAJIAGHAYI, M., LAHAIE, S., PENNOCK, D., SEDDIGHIN, M., SEDDIGHIN, S., AND YAMI, H. 2019. Fair allocation of indivisible goods to asymmetric agents. *Journal of Artificial Intelligence Research 64*, 1–20.

GAN, J., LI, B., AND WU, X. 2021. Approximately envy-free budget-feasible allocation. *CoRR abs/2106.14446*.

GARG, N., KAVITHA, T., KUMAR, A., MEHLHORN, K., AND MESTRE, J. 2010. Assigning papers to referees. *Algorithmica 58,* 1, 119–136.

GOLDBERG, P. W., HOLLENDER, A., AND SUKSOMPONG, W. 2020. Contiguous cake cutting: Hardness results and approximation algorithms. *Journal of Artificial Intelligence Research 69*, 109–141.

GOLDMAN, J. AND PROCACCIA, A. D. 2014. Spliddit: Unleashing fair division algorithms. *ACM SIGecom Exchanges 13,* 2, 41–46.

GOURVÈS, L. AND MONNOT, J. 2019. On maximin share allocations in matroids. *Theoretical Computer Science 754*, 50–64.

GOURVÈS, L., MONNOT, J., AND TLILANE, L. 2014. Near fairness in matroids. In *Proceedings of the 21st European Conference on Artificial Intelligence (ECAI)*. 393–398.

HOSSEINI, H., IGARASHI, A., AND SEARNS, A. 2020. Fair division of time: Multi-layered cake cutting. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*. 182–188.

HUMMEL, H. AND HETLAND, M. L. 2021a. Fair allocation of conflicting items. *CoRR abs/2104.06280*.

HUMMEL, H. AND HETLAND, M. L. 2021b. Guaranteeing half-maximin shares under cardinality constraints. *CoRR abs/2106.07300*.

IGARASHI, A. AND MEUNIER, F. 2021. Envy-free division of multi-layered cakes. In *Proceedings of the 17th Conference on Web and Internet Economics (WINE)*. Forthcoming.

IGARASHI, A. AND PETERS, D. 2019. Pareto-optimal allocation of indivisible goods with connectivity constraints. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*. 2045–2052.

IGARASHI, A. AND ZWICKER, W. S. 2021. Fair division of graphs and of tangled cakes. *CoRR abs/2102.08560*.

JOJIĆ, D., PANINA, G., AND ŽIVALJEVIĆ, R. 2021. Splitting necklaces, with constraints. *SIAM Journal on Discrete Mathematics 35,* 2, 1268–1286.

KYROPOULOU, M., SUKSOMPONG, W., AND VOUDOURIS, A. A. 2020. Almost envy-freeness in group resource allocation. *Theoretical Computer Science 841*, 110–123.

LI, Z. AND VETTA, A. 2021. The fair division of hereditary set systems. *ACM Transactions on Economics and Computation 9,* 2, 12:1–12:19.

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 Electronic Commerce (EC)*. 125–131.

LONC, Z. AND TRUSZCZYNSKI, M. 2020. Maximin share allocations on cycles. *Journal of Artificial Intelligence Research 69*, 613–655.

MANURANGSI, P. AND SUKSOMPONG, W. 2017. Asymptotic existence of fair divisions for groups. *Mathematical Social Sciences 89*, 100–108.

MARKAKIS, E. 2017. Approximation algorithms and hardness results for fair division. In *Trends in Computational Social Choice*, U. Endriss, Ed. AI Access, Chapter 12, 231–247.

MISRA, N., SONAR, C., VAIDYANATHAN, P. R., AND VAISH, R. 2021. Equitable division of a path. *CoRR abs/2101.09794*.

MOULIN, H. 2003. *Fair Division and Collective Welfare*. MIT Press.

PROCACCIA, A. D. 2016. Cake cutting algorithms. In *Handbook of Computational Social Choice*, F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, Eds. Cambridge University Press, Chapter 13, 311–329.

ROBERTSON, J. AND WEBB, W. 1998. *Cake-Cutting Algorithms: Be Fair if You Can*. Peters/CRC Press.

SEGAL-HALEVI, E. 2018. Fairly dividing a cake after some parts were burnt in the oven. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 1276–1284.

SEGAL-HALEVI, E. 2021. Fair multi-cake cutting. *Discrete Applied Mathematics 291*, 15–35.

SEGAL-HALEVI, E., HASSIDIM, A., AND AUMANN, Y. 2020. Envy-free division of land. *Mathematics of Operations Research 45,* 3, 896–922.

SEGAL-HALEVI, E. AND NITZAN, S. 2019. Envy-free cake-cutting among families. *Social Choice and Welfare 53,* 4, 709–740.

SEGAL-HALEVI, E., NITZAN, S., HASSIDIM, A., AND AUMANN, Y. 2017. Fair and square: Cake-cutting in two dimensions. *Journal of Mathematical Economics 70,* 8, 1–28.

SHAH, N. 2017. Spliddit: Two years of making the world fairer. *XRDS: Crossroads, The ACM Magazine for Students 24,* 1, 24–28.

STROMQUIST, W. 1980. How to cut a cake fairly. *American Mathematical Monthly 87,* 8, 640–644.

STROMQUIST, W. 2008. Envy-free cake divisions cannot be found by finite protocols. *Electronic Journal of Combinatorics 15*, #R11.

SU, F. E. 1999. Rental harmony: Sperner's lemma in fair division. *American Mathematical Monthly 106,* 10, 930–942.

SUKSOMPONG, W. 2018. Approximate maximin shares for groups of agents. *Mathematical Social Sciences 92*, 40–47.

SUKSOMPONG, W. 2019. Fairly allocating contiguous blocks of indivisible items. *Discrete Applied Mathematics 260*, 227–236.

THOMSON, W. 2016. Introduction to the theory of fair allocation. In *Handbook of Computational Social Choice*, F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, Eds. Cambridge University Press, Chapter 11, 261–283.

WALSH, T. 2020. Fair division: the computer scientist's perspective. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*. 4966–4972.

WU, X., LI, B., AND GAN, J. 2021. Budget-feasible maximum Nash social welfare allocation is almost envy-free. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*. 465–471.